arXiv:1401.2086v2 [cs.GT] 2 Jul 2015

# Actor-Critic Algorithms for Learning Nash Equilibria in $N$-player General-Sum Games

Prasad H.L. [*1], Prashanth L A [†2] and Shalabh Bhatnagar [‡3]

[1]Astrome Technologies Pvt Ltd, Bangalore, INDIA
[2]Institute for Systems Research, University of Maryland, College Park, US
[3]Department of Computer Science and Automation, Indian Institute of Science, Bangalore, INDIA

**Abstract**

We consider the problem of finding stationary Nash equilibria (NE) in a finite discounted general-sum stochastic game. We first generalize a non-linear optimization problem from Filar and Vrieze [2004] to a $N$-player setting and break down this problem into simpler sub-problems that ensure there is no Bellman error for a given state and an agent. We then provide a characterization of solution points of these sub-problems that correspond to Nash equilibria of the underlying game and for this purpose, we derive a set of necessary and sufficient SG-SP (Stochastic Game - Sub-Problem) conditions. Using these conditions, we develop two *actor-critic* algorithms: OFF-SGSP (model-based) and ON-SGSP (model-free). Both algorithms use a *critic* that estimates the value function for a fixed policy and an *actor* that performs descent in the policy space using a descent direction that avoids local minima. We establish that both algorithms converge, in self-play, to the equilibria of a certain ordinary differential equation (ODE), whose stable limit points coincide with stationary NE of the underlying general-sum stochastic game. On a single state non-generic game (see Hart and Mas-Colell [2005]) as well as on a synthetic two-player game setup with $810,000$ states, we establish that ON-SGSP consistently outperforms NashQ [Hu and Wellman, 2003] and FFQ [Littman, 2001] algorithms.

**Keywords:** General-sum discounted stochastic games, Nash equilibrium, multi-agent reinforcement learning, multi-timescale stochastic approximation.

## 1 Introduction

Traditional game theoretic developments have been for single-shot games where all agents participate and perform their preferred actions, receive rewards and the game is over. However, several emerging applications have the concept of multiple stages of action or most often the concept of *time* in them. One intermediate class of games to handle multiple stages of decision is called *repeated games*. However, repeated games do not provide for characterizing the influence of decisions made in one stage to future stages. Markov chains or Markov processes are a popular and widely applicable class of random processes which are used for modeling practical systems. Markov chains allow system designers to model *states* of a given system and then model the time behavior of the system by *connecting* those states via suitable probabilities for transition from *current state* to a *next state*. A popular extension to Markov chains which is used for modeling optimal control scenarios is the class of Markov Decision Processes (*MDPs*). Here, in a given state, an *action* is allowed to be selected from a set of available actions. Based on the choice of the action, a suitable *reward/cost* is obtained/incurred. Also, the action selected influences the probabilities of transition from one state to another. Shapley [1953] merged these concepts of MDPs

[*]prasad@astrome.co
[†]prashla@umd.edu
[‡]shalabh@csa.iisc.ernet.in

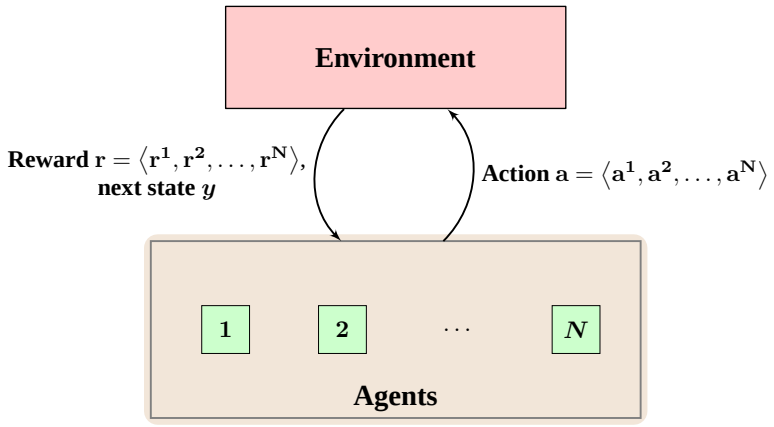


Figure 1: Multi-agent RL setting

(or basically Markov behavior) and games to come up with a new class of games called as *stochastic games*. In a stochastic game, all participating agents select their actions, each of which influence the rewards received by all agents as well as the transition probabilities. Since the inception of stochastic games by Shapley [1953], they have been an important class of models for multi-agent systems. A comprehensive treatment of stochastic games under various pay-off criteria is given by Filar and Vrieze [2004]. Many problems like fishery games, advertisement games and several oligopolistic situations can be modelled as stochastic games citep breton1986computation,filar-vrieze,olley1992dynamics,pakes1998empirical,pakes2001stochastic,bergemann1996learning.

We consider a finite *stochastic game* (also referred to as Markov game (cf. Littman [2001]) setting that evolves over discrete time instants. As illustrated in Fig. 1, at each stage and in any given state $x \in \mathcal{X}$, all agents act simultaneously with an action vector $a \in \mathcal{A}(x)$ resulting in a transition to the next state $y \in \mathcal{X}$ according to the transition probability $p(y|x,a)$ as well as a reward vector $\mathbf{r}(x,a)$. No agent gets to know what the other agents' actions are before selecting its own action and the reward $r^i(x,a)$ obtained by any agent $i$ in each stage depends on both system state $x$ (common to all agents) and the aggregate action $a$ (which includes other agents' actions). Each individual agent's *sole* objective is maximization of his/her *value function*, i.e., the expected discounted sum of rewards. However, the transition dynamics as well as the rewards depend on the actions of all agents and hence, the dynamics of the game is coupled and not independent. We assume that $\mathbf{r}(x,a)$ and the action vector $a$ are made known to all agents after every agent $i$ has picked his/her action $a^i$ - this is the canonical *model-free* setting[1]. However, we do not assume that each agent knows the other agents' policies, i.e., the distribution from which the actions are picked.

The central concept of stability in a stochastic game is that of a *Nash equilibrium*. At a Nash equilibrium point (with a corresponding Nash strategy), each agent plays a *best-response* strategy assuming all the other agents play their equilibrium strategies (see Definition 2 for a precise statement). This notion of equilibrium makes perfect sense in a game setting where agents do not have any incentive to unilaterally deviate from their Nash strategies.

Breton [1991], Filar and Vrieze [2004] establish that finding the stationary NE of a two-player discounted stochastic game is equivalent to solving an optimization problem with a non-linear objective function and linear constraints. We extend this formulation to general $N$-player stochastic games and observe that this generalization causes the constraints to be non-linear as well. Previous approaches to solving the optimization problem have not been able to guarantee convergence to global minima, even for the case of $N = 2$. In this light, our contribution is significant as we develop an algorithm to find a global minimum for any $N \geq 2$ via the following steps:[2]

**Step 1 (Sub-problems):** We break down the main optimization problem into several sub-problems. Each sub-problem can be seen as ensuring no Bellman error, for a particular state $x \in \mathcal{X}$ and agent $i \in \{1, \ldots, N\}$,

[1] While the ON-SGSP algorithm that we propose is for this setting, we also propose another algorithm - OFF-SGSP - that is model based.

[2] A preliminary version of this paper, without the proofs, was published in *AAMAS 2015* - see Prasad et al. [2015]. In comparison to the conference version, this paper includes a more detailed problem formulation, formal proofs of convergence of the two proposed algorithms, some additional experiments and a revised presentation.

where $\mathcal{X}$ is the state space and $N$ is the number of agents of the stochastic game considered.

**Step 2 (Solution points):** We provide a characterization of solution points that correspond to Nash equilibria of the underlying game. As a result, we also derive a set of necessary and sufficient conditions, henceforth referred to as SG-SP (Stochastic Game - Sub-Problem) conditions.

**Step 3 (Descent direction):** Using SG-SP conditions, we derive a descent direction that avoids local minima. This is not a steepest descent direction, but a carefully chosen descent direction specific to this optimization problem, which ensures convergence only to points of global minima that correspond to SG-SP points (and hence Nash strategies).

**Step 4 (Actor-critic algorithms):** We propose algorithms that incorporate the aforementioned descent direction to ensure convergence to stationary NE of the underlying game.

The algorithms that we propose are as follows:

**OFF-SGSP.** This is an offline, centralized and model-based scheme, i.e., it assumes that the transition structure of the underlying game is known.

**ON-SGSP.** This is an online, model-free scheme that is decentralized, i.e., learning is localized to each agent with one instance of ON-SGSP running on each agent. ON-SGSP only requires that other agents' actions and rewards are observed and not their policies, i.e., maps from states to actions.

We make the assumption that for all strategies, the resulting Markov chain is irreducible and positive recurrent. This assumption is common to the analysis of previous multi-agent RL algorithms as well (cf. Hu and Wellman [1999], Littman [2001])[3]. To the best of our knowledge, ON-SGSP is the first model-free online algorithm that converges in self-play to stationary NE for any finite discounted general-sum stochastic game where the aforementioned assumption holds.

As suggested by Bowling and Veloso [2001], two desirable properties of any multi-agent learning algorithm are as follows:

(a) *Rationality*[4]*:* Learn to play optimally when other agents follow stationary strategies; and

(b) *Self-play convergence:* Converge to a Nash equilibrium assuming all agents are using the same learning algorithm.

Our ON-SGSP algorithm can be seen to meet both the properties mentioned above. However, unlike the repeated game setting of [Bowling and Veloso, 2001, Conitzer and Sandholm, 2007], ON-SGSP solves discounted general-sum stochastic games and possesses theoretical convergence guarantees as well.

As illustrated in Fig. 2, the basic idea in both OFF-SGSP and ON-SGSP is to employ two recursions, referred to as the actor and the critic, respectively. Conceptually, these can be seen as two nested loops that operate as follows:

**Critic recursion:** This performs policy evaluation, i.e., estimates the value function for a fixed policy. In the model-based setting (i.e., of OFF-SGSP), this corresponds to the well-known dynamic programming procedure - value iteration. On the other hand, in the model-free setting (i.e., of ON-SGSP), the critic is based on temporal difference (TD) learning [Sutton and Barto, 1998].

**Actor recursion:** This incrementally updates the policy using gradient descent. For this purpose, it uses a descent direction that ensures convergence to a global minimum (and hence NE) of the optimization problem we mentioned earlier.

[3] For the case of stochastic games where there are multiple communicating classes of states or even transient states, a possible work-around is to re-start the game periodically in a random state.

[4] The term *rationality* is not to be confused with its common interpretation in economics parlance.

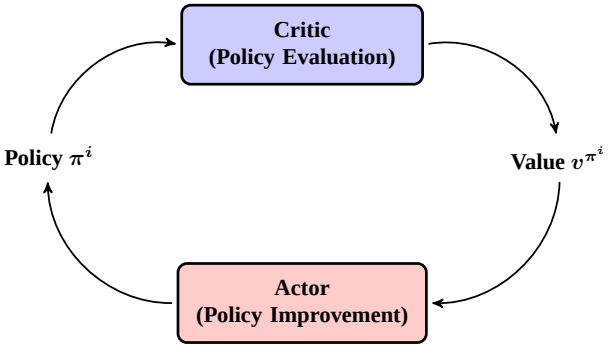


Figure 2: Operational flow of our algorithms

Using multi-timescale stochastic approximation (see Chapter 6 of Borkar [2008]) both the recursions above are run simultaneously, albeit with varying step-size parameters and this mimics the nested two-loop procedure outlined above.

The formal proof of convergence requires considerable sophistication, as we base our approach on the ordinary differential equations (ODE) method for stochastic approximation [Borkar, 2008]. While a few previous papers in the literature have adopted this approach (cf. Akchurina [2009],Weibull [1996]), their results do not usually start with an algorithm that is shown to track an ODE. Instead, an ODE is reached first via analysis and an approximate method is used to solve this ODE. On the other hand, we adopt the former approach and show that both OFF-SGSP and ON-SGSP converge using the following steps:

**1.** Using two-timescale stochastic approximation, we show that the value and policy updates on the fast and slow timescales, converge respectively to the limit points of a system of ODEs.

**2.** Next, we provide a simplified representation for the limiting set of the policy ODE and use this to establish that the asymptotically stable limit points of the policy ODE correspond to SG-SP (and hence Nash) points.

While the first step above uses a well-known result (Kushner-Clark lemma) for analysing stochastic approximation recursions, the techniques used in the second step above are quite different from those used previously. The latter step is crucial in establishing overall convergence, as the strategy $\pi$ corresponding to each stable limit gives a stationary NE of the underlying general-sum discounted stochastic game.

We demonstrate the practicality of our algorithms on two synthetic two-player setups. The first is a single state non-generic game adopted from Hart and Mas-Colell [2005] that contains two NEs (one pure, the other mixed), while the second is a *stick-together game* with $810,000$ states (to the best of our knowledge, previous works on general-sum games have never considered state spaces of this size). On the first setup, we show that ON-SGSP always converges to NE, while NashQ [Hu and Wellman, 2003] and FFQ [Littman, 2001] do not in a significant number of experimental runs. On the second setup, we show that ON-SGSP outperforms NashQ and FFQ, while exhibiting a relatively quick convergence rate - requiring approximately $21$ iterations per state.

**Map of the paper.** Section 2 reviews relevant previous works in the literature. Section 3 formalizes a general-sum stochastic game and sets the notation used throughout the paper. Section 4 formulates a general non-linear optimization problem for solving stochastic games and also form sub-problems whose solutions correspond to Nash strategies. Section 5 presents necessary and sufficient SG-SP conditions for Nash equilibria of general-sum stochastic games. Section 6 presents the offline algorithm OFF-SGSP, while Section 7 provides its online counterpart ON-SGSP. Section 8 sketches the proof of convergence for both the algorithms. Simulation results for the single state non-generic game and the stick-together game settings are presented in Section 9. Finally, concluding remarks are provided in Section 10.

## 2 Related Work

Various approaches have been proposed in the literature for computing Nash equilibrium of general-sum discounted stochastic games and we discuss some of them below.

**Multi-agent RL.** Littman [Littman, 1994] proposed a minimax Q-learning algorithm for two-player zero-sum stochastic games. Hu and Wellman [Hu and Wellman, 1999, 2003] extended the Q-learning approach to general-sum games, but their algorithms do not provide meaningful convergence guarantees. Friend-or-foe Q-learning (FFQ) [Littman, 2001] is a further improvement based on Q-learning and with guaranteed convergence. However, FFQ converges to Nash equilibria only in restricted settings (See conditions A and B in [Littman, 2001]). Moreover, the approaches in [Hu and Wellman, 1999, 2003] require computation of Nash equilibria of a bimatrix game, while the approach of Littman [2001] requires solving a linear program, in each round of their algorithms and this is a computationally expensive operation. In contrast, ON-SGSP does not require any such equilibria computations. Zinkevich et al. [2006] show that the traditional Q-learning based approaches are not sufficient to compute Nash equilibria in general-sum games[5].

**Policy hill climbing.** This is a category of previous works that is closely related to ON-SGSP algorithm that we propose. Important references here include Bowling and Veloso [2001], Bowling [2005], Conitzer and Sandholm [2007] as well as Zhang and Lesser [2010]. All these algorithms are gradient-based, model-free and are proven to converge to NE for stationary opponents in self-play. However, these convergence guarantees are for repeated games only, i.e., the setting is a single state stochastic game, where the objective is to learn the Nash strategy for a stage-game (see Definition 1 in [Conitzer and Sandholm, 2007]) that is repeatedly played. On the other hand, we consider general-sum stochastic games where the objective is to learn the best-response strategy against stationary opponents in order to maximize the value function (which is an infinite horizon discounted sum). Further, we work with a more general state space that is not restricted to be a singleton.

**Homotopy.** Herings and Peeters [2004] propose an algorithm where a homotopic path between equilibrium points of $N$ independent MDPs and the $N$ player stochastic game in question, is traced numerically. This, in turn, gives a Nash equilibrium point of the stochastic game of interest. This approach is extended by Herings and Peeters [2006] and Borkovsky et al. [2010]. OFF-SGSP shares similarities with the aforementioned homotopic algorithms in the sense that both are

**(i)** offline and model-based as they assume complete information (esp. the transition dynamics) about the game; and

**(ii)** the computational complexity for each iteration of both algorithms grows exponentially with the number of agents $N$.

**(iii)** Further, both algorithms are proven to converge to stationary NE, though their approach adopted is vastly different. OFF-SGSP is a gradient descent algorithm designed to converge to the global minimum of a nonlinear program, while the algorithm by Herings and Peeters [2004] involves a tracing procedure to find an equilibrium point.

**Linear programming.** Mac Dermed and Isbell [2009] solve stochastic games by formulating intermediate optimization problems, called Multi-Objective Linear Programs (MOLPs). However, the solution concept there is of *correlated equilibria* and Nash points are a strict subset of this class (and hence are harder to find). Also, the complexity of their algorithm scales exponentially with the problem size.

Both homotopy and linear programming methods proposed by Mac Dermed and Isbell [2009] and Herings and Peeters [2004] are tractable only for small-sized problems. The computational complexity of these algorithms may render them infeasible on large state games. In contrast, ON-SGSP is a model-free algorithm with a per-iteration complexity that is linear in $N$, allowing for practical implementations on large state game settings (see Section 9 for one such example with a state space cardinality of $810,000$). We however mention that per-iteration complexity alone is not sufficient to quantify the performance of an algorithm - see Remark 7.

[5] We avoid this impossibility result by searching for both values and policies instead of just values, in our proposed algorithms.

**Rational learning.** A popular algorithm with guaranteed convergence to Nash equilibria in general-sum stochastic games is rational learning, proposed by Kalai and Lehrer [1993]. In their algorithm, each agent $i$ maintains a prior on what he believes to be other agents' strategy and updates it in a Bayesian manner. Combining this with certain assumptions of absolute continuity and grain of truth, the algorithm there is shown to converge to NE. ON-SGSP operates in a similar setting as that in [Kalai and Lehrer, 1993], except that we do not assume the knowledge of reward functions. ON-SGSP is a model-free online algorithm and unlike [Kalai and Lehrer, 1993], any agent's strategy in ON-SGSP does not depend upon Bayesian estimates of other agents' strategies and hence, their absolute continuity/grain of truth assumptions do not apply.

**Evolutionary algorithm.** Akchurina [2009] employs numerical methods in order to solve a system of ODEs and only establishes empirical convergence to NE for a group of randomly generated games. In contrast, ON-SGSP is a model-free algorithm that is provably convergent to NE in self-play. We also note that the system of ODEs given by Akchurina [2009] (also found in [Weibull, 1996, pp. 189]) turns out to be similar to a portion of the ODEs that are tracked by ON-SGSP.

**Remark 1.** *Shoham et al. [2003] and Shoham et al. [2007] question if Nash equilibrium is a useful solution concept for general-sum games. However, if we are willing to concede that* prescriptive, equilibrium agenda *is indeed useful for stochastic games, then we believe our work is theoretically significant. Our ON-SGSP algorithm is a prescriptive, co-operative learning algorithm that observes a sample path from the underlying game and converges to stationary NE. To the best of our knowledge, this is the first algorithm to do so, with proven convergence.*

# 3 Formal Definitions

A stochastic game can be seen to be an extension of the single-agent Markov decision process. A discounted reward stochastic game is described by a tuple $< N, \mathcal{X}, \mathcal{A}, p, \mathbf{r}, \beta >$, where $N$ represents the number of agents, $\mathcal{X}$ denotes the state space and $\mathcal{A} = \cup_{x \in \mathcal{X}} \mathcal{A}(x)$ is the aggregate action space, where $\mathcal{A}(x) = \prod_{i=1}^{N} \mathcal{A}^i(x)$ is the Cartesian product of action spaces $(\mathcal{A}^i(x))$ of individual agents when the state of the game is $x \in \mathcal{X}$. We assume both state and action spaces to be finite. Let $p(y|x,a)$ denote the probability of going from the current state $x \in \mathcal{X}$ to $y \in \mathcal{X}$ when the vector of actions $a \in \mathcal{A}(x)$ (of the $N$ players) is chosen and let $\mathbf{r}(x,a) = \langle r^i(x,a) : i = 1, 2, \ldots, N \rangle$ denote the vector of reward functions of all agents when the state is $x \in \mathcal{X}$ and the vector of actions $a \in \mathcal{A}(x)$ is chosen. Also, $0 < \beta < 1$ denotes the discount factor that controls the influence of the rewards obtained in the future on the agents' strategy (see Definition 1 below).

**Notation.** $\langle \cdots \rangle$ represents a column vector and $\underline{1}_m$ is a vector of ones with $m$ elements. The various constituents of the stochastic game considered are denoted as follows:[6]

**Action:** $a = \langle a^1, a^2, \ldots, a^N \rangle \in \mathcal{A}(x)$ is the aggregate action, $a^{-i}$ is the tuple of actions of all agents except $i$ and $\mathcal{A}^{-i}(x) := \prod_{j \neq i} \mathcal{A}^j(x)$ is the set of feasible actions in state $x \in \mathcal{X}$ of all agents except $i$.

**Policy:** $\pi^i(x, a^i)$ is the probability of picking action $a^i \in \mathcal{A}^i(x)$ by agent $i$ in state $x \in \mathcal{X}$,
$\pi^i(x) = \langle \pi^i(x, a^i) : a^i \in \mathcal{A}^i(x) \rangle$ is the randomized policy vector in state $x \in \mathcal{X}$ for the agent $i$, $\pi^i = \langle \pi^i(x) : x \in \mathcal{X} \rangle$, $\pi = \langle \pi^i : i = 1, 2, \ldots, N \rangle$ is the strategy-tuple of all agents and
$\pi^{-i} = \langle \pi^j : j = 1, 2, \ldots, N, j \neq i \rangle$ is the strategy-tuple of all agents except agent $i$. We focus only on *stationary strategies* in this paper, as suggested by Theorem 1.

**Transition Probability:** Let $\pi(x,a) = \prod_{i=1}^{N} \pi^i(x, a^i)$ and $\pi^{-i}(x, a^{-i}) = \prod_{j=1, j \neq i}^{N} \pi^j(x, a^j)$. Then, the (Markovian) transition probability from state $x \in \mathcal{X}$ to state $y \in \mathcal{X}$ when each agent $i$ plays according to its randomized

[6] We use the terms *policy* and *strategy* interchangeably in the paper.

strategy $\pi^i$ can be written as:

$$p(y|x,\pi) = \sum_{a \in \mathcal{A}(x)} p(y|x,a)\pi(x,a).$$

**Reward:** $r^i(x,a)$ is the single-stage reward obtained by agent $i$ in state $x \in \mathcal{X}$, where $a \in \mathcal{A}(x)$ is the aggregate action taken.

**Definition 1.** ***(Value function)*** *The value function is the expected return for any agent $i \in \{1,2,\ldots,N\}$ and is defined as*

$$v^i_\pi(s_0) = \mathbb{E}\left[\sum_t \beta^t \sum_{a \in \mathcal{A}(x)} \left(r^i(s_t,a)\pi(s_t,a)\right)\right]. \tag{1}$$

Given the above notion of the value function, the goal of each agent is to find a strategy that achieves a *Nash equilibrium*. The latter is defined as follows:

**Definition 2.** ***(Nash Equilibrium)*** *A stationary Markov strategy $\pi^* = \langle \pi^{1*}, \pi^{2*}, \ldots, \pi^{N*} \rangle$ is said to be Nash if*

$$v^i_{\pi^*}(s) \geq v^i_{\langle \pi^i, \pi^{-i*} \rangle}(s), \forall \pi^i, \forall i, \forall s \in \mathcal{X}.$$

*The corresponding equilibrium of the game is said to be Nash equilibrium.*

Since we consider a discounted stochastic game with a finite state space, we have the following well-known result that ensures the existence of stationary equilibrium:

**Theorem 1.** *Any finite discounted stochastic game has an equilibrium in stationary strategies.*

We shall refer to such stationary randomized strategies as *Nash strategies*. The reader is referred to Fink [1964], Takahashi [1964], Sobel [1971] for a proof of Theorem 1.

# 4 A Generalized Optimization Problem

**Basic idea.** Using dynamic programming the Nash equilibrium condition in Definition 2 can be re-written as: $\forall x \in \mathcal{X}, \forall i = 1,2,\ldots,N$,

$$v^i_{\pi^*}(x) = \max_{\pi^i(x) \in \Delta(\mathcal{A}^i(x))} \left\{ E_{\pi^i(x)} Q^i_{\pi^{-i*}}(x,a^i) \right\}, \tag{2}$$

where

$$Q^i_{\pi^{-i}}(x,a^i) = E_{\pi^{-i}(x)}\left[ r^i(x,a) + \beta \sum_{y \in U(x)} p(y|x,a) v^i(y) \right],$$

represents the marginal value associated with picking action $a^i \in \mathcal{A}^i(x)$, in state $x \in \mathcal{X}$ for agent $i$, while other agents act according to $\pi^{-i}$. Also, $\Delta(\mathcal{A}^i(x))$ denotes the set of all possible probability distributions over $\mathcal{A}^i(x)$.

The basic idea behind the optimization problem that we formulate below is to model the objective such that the value function is correct w.r.t. the agents' strategies, while adding a constraint to ensure that a feasible solution to the problem corresponds to Nash equilibrium.

**Objective.** A possible optimization objective for agent $i$ would be

$$f^i(\mathbf{v}^i, \pi) = \sum_{x \in \mathcal{X}} \left( v^i(x) - E_{\pi^i} Q^i_{\pi^{-i}}(x,a^i) \right).$$

The objective above has to be minimized over all policies $\pi^i \in \Delta(\mathcal{A}^i(x))$. But $Q^i_{\pi^{-i}}(x,a^i)$, by definition, is dependent on strategies of all other agents. So, an isolated minimization of $f^i(\mathbf{v}^i, \pi^i)$ would not be meaningful

and hence, we consider the aggregate objective $f(\mathbf{v}, \pi) = \sum_{i=1}^{N} f^i(\mathbf{v}^i, \pi)$. This objective is minimized over all policies $\pi^i \in \Delta(\mathcal{A}^i(x))$ of all agents. Thus, we have an optimization problem with objective as $f(\mathbf{v}, \pi)$ along with the natural constraints ensuring that the policy vectors $\pi^i(x)$ remain as probabilities over all feasible actions $\mathcal{A}^i(x)$ in all states $x \in \mathcal{X}$ and for agents $i = 1, \ldots, N$.

**Constraints.** Notice that an optimization problem with the objective discussed above has only a set of simple constraints ensuring that $\pi$ remains a valid strategy. However, this is not sufficient to accurately represent Nash equilibria of the underlying game. Here, we look at a possible set of additional constraints which might make the optimization problem more useful. Note that the term being maximized in equation (2), i.e., $E_{\pi^i} Q^i_{\pi^{-i}}(x, a^i)$, represents a convex combination of the values of $Q^i_{\pi^{-i}}(x, a^i)$ over all possible actions $a^i \in \mathcal{A}^i(x)$ in a given state $x \in \mathcal{X}$ for a given agent $i$. Thus, it is implicitly implied that

$$Q^i_{\pi^{-i}}(x, a^i) \leq v^i_{\pi^*}(x), \forall a^i \in \mathcal{A}^i(x), x \in \mathcal{X}, i = 1, 2, \ldots, N.$$

Formally, the optimization problem for any $N \geq 2$ is given below:

$$\left.\begin{array}{l} \min\limits_{\mathbf{v},\pi} f(\mathbf{v}, \pi) = \sum\limits_{i=1}^{N} \sum\limits_{x \in \mathcal{X}} \left( v^i(x) - E_{\pi^i} Q^i_{\pi^{-i}}(x, a^i) \right) \text{s.t.} \\ (a)\, \pi^i(x, a^i) \geq 0, \forall a^i \in \mathcal{A}^i(x), x \in \mathcal{X}, i = 1, 2, \ldots, N, \\ (b)\, \sum\limits_{i=1}^{N} \pi^i(x, a^i) = 1, \forall x \in \mathcal{X}, i = 1, 2, \ldots, N, \\ (c)\, Q^i_{\pi^{-i}}(x, a^i) \leq v^i(x), \forall a^i \in \mathcal{A}^i(x), x \in \mathcal{X}, i = 1, 2, \ldots, N. \end{array}\right\} \quad (3)$$

In the above, 3(a)–3(b) ensure that $\pi$ is a valid policy, while 3(c) is necessary for any valid policy to be a NE of the underlying game.

**Theorem 2.** *A feasible point $(v^*, \pi^*)$ of the optimization problem (3) provides a Nash equilibrium in stationary strategies to the corresponding general-sum discounted stochastic game if and only if $f(v^*, \pi^*) = 0$.*

*Proof.* See [Filar and Vrieze, 2004, Theorem 3.8.2] for a proof in the case of $N = 2$. The proof works in a similar manner for general $N$. □

A few remarks about the difficulties involved in solving (3) are in order.

**Remark 2.** ***(Non-linearity)*** *For the case of $N = 2$, the objective $f(\mathbf{v}, \pi)$ in* (3) *is of order* 3*, while the toughest constraint 3(c) is quadratic. This is apparent from the fact that the second term inside the summation in $f(\mathbf{v}, \pi)$ has the following variables multiplied: $\pi_1$ in the first expectation, $\pi_2$ inside the expectation in the definition of the $Q$-function and $v$ inside the second term of the expectation in $Q$-function. Along similar lines, the constraint 3(c) can be inferred to be quadratic. Thus, we have an optimization problem with a third-order objective and quadratic constraints, even for the case of $N = 2$ and the constituent functions (both objective and constraints) can be easily seen to be non-linear. For a general $N > 2$, the problem* (3) *gets more complicated, as more policy variables $\pi_1, \ldots, \pi_N$ are thrown in.*

**Remark 3.** ***(Beyond local minima)*** *Filar and Vrieze [2004] have formulated a non-linear optimization problem for the case of two-player* zero-sum *stochastic games. An associated result (Theorem 3.9.4, page 141, in [Filar and Vrieze, 2004]) states that every local minimum of that optimization problem is also a global minimum. However, this result does not hold for a general-sum game even in the two-player (and also $N \geq 3$) setting and hence, the requirement is for a global optimization scheme that solves* (3)*.*

**Remark 4.** ***(No steepest descent)*** *From the previous remark, it is apparent that a simple* steepest descent *scheme is not enough to solve* (3) *even for the two-player setting. This is because there can be local minima of* (3) *that do not correspond to the global minimum and steepest descent schemes guarantee convergence to local minima only. Note that steepest descent schemes were sufficient to solve for Nash equilibrium strategies in two-player zero-sum*

*stochastic games, while this is not the case with general-sum $N$-player games, with $N \geq 2$. Sections 3 and 4 of [Prasad and Bhatnagar, 2015] contain a detailed discussion on inadequacies of steepest descent for general-sum games.*

**Remark 5.** ***(No Newton method)*** *A natural question that arises is can one employ a Newton method in order to solve* (3) *and the answer is in the negative. Observe that the Hessian of the objective function $f(\mathbf{v}, \pi)$ in* (3) *has its diagonal elements to be zero and this makes it indefinite. This make Newton methods infeasible as they require invertibility of the Hessian to work.*

We overcome the above difficulties by deriving a descent direction (that is not necessarily steepest) in order to solve (3). Before we present the descent direction, we break-down (3) into simpler sub-problems. Subsequently we descibe Stochastic Game - Sub-Problem (SG-SP) conditions that are both necessary and sufficient for the problem (3).

### Sub-problems for each state and agent

We form sub-problems from the main optimization problem (3) along the lines of Prasad and Bhatnagar [2012], for each state $x \in \mathcal{X}$ and each agent $i \in \{1, 2, \ldots, N\}$. The sub-problems are formed with the objective of ensuring that there is no Bellman error (see $g^i_{x,z}(\theta)$ below). For any $x \in \mathcal{X}$, $z = 1, 2, \ldots, |\mathcal{A}^i(x)|$ and $i \in \{1, 2, \ldots, N\}$, let $\theta := \left\langle \mathbf{v}^i, \pi^{-i}(x) \right\rangle$ denote the value-policy tuple and let

$$g^i_{x,z}(\theta) := Q^i_{\pi^{-i}}(x, a^i_z) - v^i(x) \tag{4}$$

denote the Bellman error. Further, let $p_z := \pi^i(x, a^i_z)$ and $p = \left\langle p_z : z = 1, 2, \ldots, |\mathcal{A}^i(x)| \right\rangle$. Then, the sub-problems are formulated as follows:

$$\begin{aligned} &\min_{\theta, p} h_x(\theta, p) = \sum_{z=1}^{|\mathcal{A}^i(x)|} p_z \left[ -g^i_{x,z}(\theta) \right] \\ &\text{s.t. } g^i_{x,z}(\theta) \leq 0, -p_z \leq 0, \text{ for } z = 1, 2, \ldots, |\mathcal{A}^i(x)|, \\ &\text{and } \sum_z p_z = 1. \end{aligned} \tag{5}$$

# 5 Stochastic Game - Sub-Problem (SG-SP) Conditions

In this section, we derive a set of necessary and sufficient conditions for solutions of (3) and establish their equivalence with Nash strategies.

**Definition 3** (SG-SP Point)**.** *A point $(\mathbf{v}^*, \pi^*)$ of the optimization problem (3) is said to be an SG-SP point if it is a feasible point of (3) and for every sub-problem, i.e., for all $x \in \mathcal{X}$ and $i \in \{1, 2, \ldots, N\}$,*

$$p^*_z g^i_{x,z}(\theta^*) = 0, \qquad \forall z = 1, 2, \ldots, |\mathcal{A}^i(x)|. \tag{6}$$

The above conditions, which define a point to be an SG-SP point, are called SG-SP conditions.

## 5.1 Equivalence of SG-SP with Nash strategies

The connection between SG-SP points and Nash equilibria can be seen intuitively as follows:
**(i)** The objective function $f(v^*, \pi^*)$ in (3) can be expressed as a summation of terms of the form $p^*_z[-g^i_{x,z}(\theta^*)]$ over $z = 1, 2, \ldots, |\mathcal{A}^i(x)|$ and over all sub-problems. Condition (6) suggests that each of these terms is zero which implies $f(v^*, \pi^*) = 0$.
**(ii)** The objective of the sub-problem is to ensure that there is no Bellman error, which in turn implies that the value estimates $v^*$ are correct with respect to the policy $\pi^*$ of all agents.
Combining (i) and (ii) with Theorem 3.8.2 of [Filar and Vrieze, 2004], we have the following result:

**Theorem 3** (Nash $\Leftrightarrow$ SG-SP)**.** *A strategy $\pi^*$ is Nash if and only if $(v^*, \pi^*)$ for the corresponding optimization problem (3) is an SG-SP point.*

The proof of the above theorem follows from a combination of Lemmas 1 and 2, presented below.

**Lemma 1** (SG-SP $\Rightarrow$ Nash)**.** *An SG-SP point $(\mathbf{v}^*, \pi^*)$ gives Nash strategy-tuple for the underlying stochastic game.*

*Proof.* The objective function value $f(v^*, \pi^*)$ of the optimization problem (3) can be expressed as a summation of terms of the form $p_z^*[-g_{x,z}^i(\theta^*)]$ over $z = 1, 2, \ldots, m$ and over all sub-problems. Condition (6) suggests that each of these terms is zero which implies $f(v^*, \pi^*) = 0$. From Filar and Vrieze [2004, Theorem 3.8.2, page 132], since $(v^*, \pi^*)$ is a feasible point of (3) and $f(v^*, \pi^*) = 0$, $(v^*, \pi^*)$ corresponds to Nash equilibrium of the underlying stochastic game. □

**Lemma 2** (Nash $\Rightarrow$ SG-SP)**.** *A strategy $\pi^*$ is Nash if $(v^*, \pi^*)$ for the corresponding optimization problem (3) is an SG-SP point.*

*Proof.* From Filar and Vrieze [2004, Theorem 3.8.2, page 132], if a strategy $\pi^*$ is Nash, then a feasible point $(v^*, \pi^*)$ exists for the corresponding optimization problem (3), where $f(v^*, \pi^*) = 0$. From the constraints of (3), it is clear that for a feasible point, $p_z^*[-g_{x,z}^i(\theta^*)] \geq 0$, for $z = 1, 2, \ldots, m$, for every sub-problem. Since the sum of all these terms, i.e., $f(v^*, \pi^*)$, is zero, each of these terms is zero, i.e., $(v^*, \pi^*)$ satisfies (6). Thus, $(v^*, \pi^*)$ is an SG-SP point. □

## 5.2 Kinship to Karush-Kuhn-Tucker - Sub-Problem (KKT-SP) conditions

Prasad and Bhatnagar [2012] consider a similar optimization problem as (3) for the case of two agents, i.e., $N = 2$ and derive a set of verifiable necessary and sufficient conditions that they call KKT-SP conditions. In the following, we first extend the KKT-SP conditions to $N$-player stochastic games, for any $N \geq 2$ and later establish the equivalence between KKT-SP conditions with the SG-SP conditions formulated above.

The Lagrangian corresponding to (5) can be written as

$$k(\theta, p, \lambda, \delta, s, t) = h_x(\theta, p) + \sum_{z=1}^{|\mathcal{A}^i(x)|} \lambda_z \left( g_{x,z}^i(\theta) + s_z^2 \right) + \sum_{z=1}^{|\mathcal{A}^i(x)|} \delta_z \left( -p_z + t_z^2 \right), \tag{7}$$

where $\lambda_z$ and $\delta_z$ are the Lagrange multipliers and $s_z$ and $t_z, z = 1, 2, \ldots, |\mathcal{A}^i(x)|$ are the slack variables, corresponding to the first and second constraints of the sub-problem (5), respectively.

Using the Lagrangian (7), the associated KKT conditions for the sub-problem (5) corresponding to a state $x \in S$ and agent $i \in \{1, \ldots, N\}$ at a point $\langle \theta^*, p^* \rangle$ are the following:

$$\left.\begin{aligned}
&(a)\, \nabla_\theta h_x(\theta^*, p^*) + \sum_{z=1}^{m} \lambda_z \nabla_\theta g_{x,z}^i(\theta^*) = 0, \\
&(b)\, \frac{\partial h_x(\theta^*, p^*)}{\partial p_z} - \delta_z + \delta_m = 0, \qquad z = 1, 2, \ldots, m, \\
&(c)\, \delta_z p_z^* = 0, \qquad z = 1, 2, \ldots, m, \\
&(d)\, \lambda_z g_{x,z}^i(\theta^*) = 0, \qquad z = 1, 2, \ldots, m, \\
&(e)\, \lambda_z \geq 0, \qquad z = 1, 2, \ldots, m, \\
&(f)\, \delta_z \geq 0, \qquad z = 1, 2, \ldots, m.
\end{aligned}\right\} \tag{8}$$

KKT-SP conditions are shown to be necessary and sufficient for $(\mathbf{v}^*, \pi^*)$ to represent a Nash equilibrium point of the underlying stochastic game and $\pi^*$ to be a Nash strategy-tuple in the case of $N = 2$ (see Theorem 3.8 in [Prasad and Bhatnagar, 2012]). However, this requires an additional assumption that for each sub-problem,

$\left\{\nabla_\theta g^i_{x,z}(\theta^*) : z = 1, 2, \ldots, m\right\}$ is a set of linearly independent vectors. On the other hand, the SG-SP conditions (see Definition 3) that we formulate do not impose any additional linear independence requirement, in order to ensure that the solution points of the sub-problems correspond to Nash equilibria.

The following lemma establishes the kinship between SG-SP and KKT-SP conditions.

**Lemma 3** (KKT-SP $\Rightarrow$ SG-SP)**.** *A KKT-SP point is also an SG-SP point.*

*Proof.* A KKT-SP point $(\mathbf{v}^*, \pi^*)$ is a feasible point of (3). For every sub-problem, substitute and eliminate $\lambda^*_z = p^*_z$ and $\delta^*_z = -g^i_{x,z}(\theta^*)$, $z = 1, 2, \ldots, m$. Then

1. Conditions (8(a)) and (8(b)) are satisfied;
2. Conditions (8(c)) and (8(d)) reduce to (6); and
3. Conditions (8(e)) and (8(f)) are satisfied as the point $(\mathbf{v}^*, \pi^*)$ is assumed to be feasible.

□

From the above, it is evident that the simpler and more general (for any $N$) SG-SP conditions can be used for Nash equilibria as compared to KKT-SP conditions because:
**(i)** every KKT-SP point is also an SG-SP point and
**(ii)** SG-SP conditions do not impose any additional linear independence requirement in order to be Nash points.

# 6 OFF-SGSP: Offline, Model-Based

**Basic idea.** As outlined in the introduction, OFF-SGSP is an actor-critic algorithm that operates using two timescale recursions as follows

**Critic recursion:** This estimates the value function $\mathbf{v}$ using value iteration, along the faster timescale; and

**Actor recursion:** This operates along the slower timescale and updates the policy in the descent direction so as to ensure convergence to an SG-SP point.

As mentioned before, OFF-SGSP is a model-based algorithm and the transition dynamics and reward structure of the underlying game are used for both steps above.

**Update rule.** Using two timescale recursions, OFF-SGSP updates the value-policy tuple $(v, \pi)$ as follows: For all $x \in \mathcal{X}$ and $a^i \in \mathcal{A}^i(x)$,

$$\textbf{Actor: } \pi^i_{n+1}(x, a^i) = \Gamma\left(\pi^i_n(x, a^i) - b(n)\sqrt{\pi^i_n(x, a^i)}\left|g^i_{x,a^i}(\mathbf{v}^i_n, \pi^{-i}_n)\right| \overline{\text{sgn}}\left(\frac{\partial f(\mathbf{v}_n, \pi_n)}{\partial \pi^i}\right)\right), \tag{9}$$

$$\textbf{Critic: } v^i_{n+1}(x) = v^i_n(x) + c(n) \sum_{a^i \in \mathcal{A}^i(x)} \pi^i_n(x, a^i) g^i_{x,a^i}(\mathbf{v}^i_n, \pi^{-i}_n), \tag{10}$$

where $g^i_{x,a^i}(\mathbf{v}^i, \pi^{-i}) := Q^i_{\pi^{-i}}(x, a^i) - v^i(x)$ denotes the Bellman error, $f(\mathbf{v}, \pi)$ is the objective function in (3) and $\Gamma$ is a projection operator that ensures that the updates to $\pi$ stay within the simplex $\mathcal{D} = \{(d_2, \ldots, d_{|\mathcal{A}^i(x)|}) \mid d_i \geq 0, \forall i = 1, \ldots, |\mathcal{A}^i(x)|, \sum_{j=2}^{|\mathcal{A}^i(x)|} d_j \leq 1\}$. Further, using $\Gamma$, one ensures that $d_1 = 1 - \sum_{j=2}^{|\mathcal{A}^i(x)|}, d_j \in [0, 1]$. Here $\overline{\text{sgn}}(\cdot)$ is a continuous version of the sign function and projects any $x$, outside of a very small interval around

$0$, to $\pm 1$ according to the sign of $x$ (see Remark 10 for a precise definition). Continuity of $\overline{\text{sgn}}(\cdot)$ is a technical requirement that helps in providing strong convergence guarantees[7].

The following assumption on the step-sizes ensures that the $\pi$-recursion (9) proceeds on a slower timescale in comparison to the $v$-recursion (10):

**Assumption 1.** *The step-sizes $\{b(n)\}, \{c(n)\}$ satisfy*

$$\sum_{n=1}^{\infty} b(n) = \sum_{n=1}^{\infty} c(n) = \infty, \sum_{n=1}^{\infty} \left(b^2(n) + c^2(n)\right) < \infty, \frac{b(n)}{c(n)} \to 0.$$

**Justification for descent direction.** The following proposition proves that the decrement for the policy in (9) is a valid descent direction for the objective function $f(\cdot, \cdot)$ in (3).

**Proposition 1.** *For each $i = 1, 2, \ldots, N, x \in \mathcal{X}, a^i \in \mathcal{A}^i(x)$, we have that $-\sqrt{\pi^i(x, a^i)} \left| g^i_{x,a^i}(\mathbf{v}^i, \pi^{-i}) \right| \overline{\text{sgn}} \left( \frac{\partial f(\mathbf{v}, \pi)}{\partial \pi^i} \right)$ is a non-ascent, and in particular a descent direction if $\sqrt{\pi^i(x, a^i)} \left| g^i_{x,a^i}(\mathbf{v}^i, \pi^{-i}) \right| \neq 0$, in the objective $f(\mathbf{v}, \pi)$ of (3).*

*Proof.* The objective $f(\mathbf{v}, \pi)$ can be rewritten as

$$f(\mathbf{v}, \pi) = \sum_{i=1}^{N} \sum_{x \in \mathcal{X}} \sum_{a^i \in \mathcal{A}^i(x)} \left\{ \pi^i(x, a^i) \left[ -g^i_{x,a^i}(\mathbf{v}^i, \pi^{-i}) \right] \right\}.$$

Assume $f(\mathbf{v}, \pi) > 0$, otherwise the solution to (3) is already achieved. For an $a^i \in \mathcal{A}^i(x)$ for some $x \in \mathcal{X}$ and $i \in \{1, 2, \ldots, N\}$, let

$$\hat{\pi}^i(x, a^i) = \pi^i(x, a^i) - \delta \sqrt{\pi^i(x, a^i)} \left| g^i_{x,a^i}(\mathbf{v}^i, \pi^{-i}) \right| \overline{\text{sgn}} \left( \frac{\partial f(\mathbf{v}, \pi)}{\partial \pi^i} \right),$$

for a small $\delta > 0$. Let $\hat{\pi}$ be the same as $\pi$ except that action $a^i$ is picked as defined above. Then by a Taylor series expansion of $f(\mathbf{v}, \hat{\pi})$ till the first order term, we obtain

$$f(\mathbf{v}, \hat{\pi}) = f(v, \pi) + \delta \left[ -\sqrt{\pi^i(x, a^i)} \left| g^i_{x,a^i}(\mathbf{v}^i, \pi^{-i}) \right| \right] \overline{\text{sgn}} \left( \frac{\partial f(\mathbf{v}, \pi)}{\partial \pi^i} \right) \frac{\partial f(\mathbf{v}, \pi)}{\partial \pi^i(x, a^i)} + o(\delta).$$

The rest of the proof amounts to showing that the second term in the expansion above is $\leq 0$. This can be inferred as follows:

$$\begin{aligned} & - \sqrt{\pi^i(x, a^i)} \left| g^i_{x,a^i}(\mathbf{v}^i, \pi^{-i}) \right| \overline{\text{sgn}} \left( \frac{\partial f(\mathbf{v}, \pi)}{\partial \pi^i} \right) \frac{\partial f(\mathbf{v}, \pi)}{\partial \pi^i} \\ = & - \sqrt{\pi^i(x, a^i)} \left| g^i_{x,a^i}(\mathbf{v}^i, \pi^{-i}) \right| \left| \frac{\partial f(\mathbf{v}, \pi)}{\partial \pi^i} \right| \leq 0, \end{aligned}$$

and is in particular $< 0$ if $\sqrt{\pi^i(x, a^i)} \left| g^i_{x,a^i}(\mathbf{v}^i, \pi^{-i}) \right| \neq 0$.

Thus, for $a^i \in \mathcal{A}^i(x)$, $x \in \mathcal{X}$ and $i \in \{1, 2, \ldots, N\}$ where $\pi^i(x, a^i) > 0$ and $g^i_{x,a^i}(\mathbf{v}^i, \pi^{-i}) \neq 0$, $f(\mathbf{v}, \hat{\pi}) < f(\mathbf{v}, \pi)$ for small enough $\delta$. The claim follows. □

[7] Using the normal $\text{sgn}()$ function is problematic for an ODE approach based analysis, as $\text{sgn}()$ is discontinuous. In other words, with $\text{sgn}()$ the underlying system for the policy $\pi^i_n$ will be a stochastic recursive inclusion and providing meaningful guarantees for such inclusions would require more assumptions. In comparison, the results are stronger for the ODE approach. This is the motivation behind employing $\overline{\text{sgn}}()$, which is a continuous extension of $\text{sgn}()$. The function $\overline{\text{sgn}}(x)$ projects any $x$ outside of a small interval around 0 (say $[-\nu, \nu]$ for some $\nu > 0$ small) to either $+1$ or $-1$ as $\text{sgn}()$ would do and within the interval $[-\nu, \nu]$, one may choose $\overline{\text{sgn}}(x) = x$ or any other continuous function with compatible end-point values. One could choose $\nu$ arbitrarily close to 0, making $\overline{\text{sgn}}$ practically very close to $sgn$.

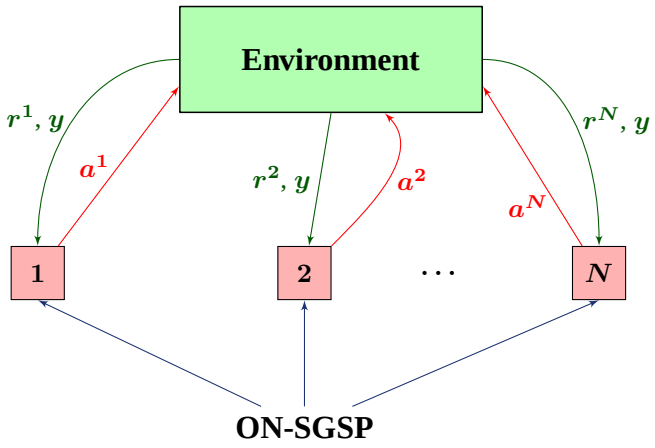


Figure 3: ON-SGSP's decentralized on-line learning model with $N$ agents

**Main result.** Let $R^i_\pi = \langle r^i(x,\pi), x \in \mathcal{X} \rangle$ be a column vector of rewards to agent $i$ and $P_\pi = [p(y|x,\pi), x \in \mathcal{X}, y \in \mathcal{X}]$ be the transition probability matrix, both for a given $\pi$. Then, the value function for a given policy $\pi$ is defined as

$$\mathbf{v}^i_\pi = [I - \beta P_\pi]^{-1} R^i_\pi, i = 1, 2, \ldots, N. \tag{11}$$

The above will be used to characterize the limit of the critic-recursion (10). Before presenting the main result, we specify the ODE that underlies the actor-recursion (9): For all $a^i \in \mathcal{A}^i(x), x \in \mathcal{X}, i = 1, 2, \ldots, N$,

$$\frac{d\pi^i(x,a^i)}{dt} = \bar{\Gamma}\left(-\sqrt{\pi^i(x,a^i)}\left|g^i_{x,a^i}(\mathbf{v}^i_\pi, \pi^{-i})\right| \overline{\text{sgn}}\left(\frac{\partial f(\mathbf{v}_\pi,\pi)}{\partial \pi^i}\right)\right), \tag{12}$$

where $\bar{\Gamma}$ is a projection operator that restricts the evolution of the above ODE to the simplex $\mathcal{D}$ (see Section 8 for a precise definition).

The main result regarding the convergence of OFF-SGSP is as follows:

**Theorem 4.** *Let $G$ denote the set of all feasible points of the optimization problem* (3) *and $K$ the set of limit points of the ODE* (12)*. Further, let $K_1 = K \cap G$ and $\mathcal{K}^* = \{(\mathbf{v}^i_{\pi^*}, \pi^*) \mid \pi^* \in K_1\}$. Then, for any agent $i = 1, \ldots, N$, the sequence of iterates $(v^i_n, \pi^i_n), n \geq 0$ satisfy*

$$(v^i_n, \pi^i_n) \to \mathcal{K}^* \text{ a.s.}$$

*Proof.* See Section 8.1. □

From the above theorem, we can infer the following:

**(i)** The set of infeasible limit points of the ODE (12), i.e., $K_2 = K \setminus K_1$, are asymptotically unstable (see Lemma 5 in Section 8 for a formal proof); and

**(ii)** OFF-SGSP converges to the set $\mathcal{K}^*$, which is the set of all asymptotically stable limit points of the system of ODEs (19). Further, $\mathcal{K}^*$ corresponds to SG-SP (and hence Nash) points and hence, OFF-SGSP is shown to converge almost surely to a NE of the underlying discounted stochastic game.

# 7 ON-SGSP: Online and Model-Free

Though OFF-SGSP is suitable for only off-line learning of Nash strategies, it is amenable for extension to the general (on-line) multi-agent RL setting where neither the transition probability $p$ nor the reward function $r$ are

explicitly known. ON-SGSP operates in the latter model-free setting and uses the stochastic game as a generative model.

As illustrated in Fig. 3, every iteration in ON-SGSP represents a discrete-time interaction with the environment, where each agent presents its action to the environment and observes the next state and the reward vector of all agents. The learning is localized to each agent $i \in \{1, 2, \ldots, N\}$, making the setting decentralized. This is in the spirit of earlier multi-agent RL approaches (cf. Hu and Wellman [1999], Hu and Wellman [2003] and Littman [2001]).

Algorithm 1 presents the complete structure of ON-SGSP along with update rules for the value and policy parameters. The algorithm operates along two timescales as follows:

**Critic (faster timescale):** Each agent estimates its own value function as well as that of other agents, using a temporal-difference (TD) [Sutton, 1988] type update in (13). Moreover, the gradient $\frac{\partial f(\mathbf{v}_n, \pi_n)}{\partial \pi^i(x, a^i)}$ is also estimated in an online manner via the $\xi$-recursion in (14). Note that the $\xi$-recursion is made necessary due to the fact that ON-SGSP operates in a model-free setting.

**Actor (slower timescale):** The policy update is similar to OFF-SGSP, except that the estimates of value $v$ and gradient $\xi$ are used to derive the decrement in (15).

Note that, since ON-SGSP operates in a model-free setting, both the value and policy updates are different in comparison to OFF-SGSP. The value **v** update (13) on the faster timescale can be seen to be the stochastic approximation variant of value iteration and it converges to the same limit as in OFF-SGSP, without knowing the model. On the other hand, the policy update (15) on the slower timescale involves a decrement that is motivated by the descent direction suggested by Proposition 1.

**Algorithm 1** ON-SGSP

**Input:** Starting state $x_0$, initial point $\theta_0^i = (\mathbf{v}_0^i, \pi_0^i)$, step-sizes $\{b(n), c(n)\}_{n\geq 1}$, number of iterations to run $M >> 0$.
**Initialization:** $n \leftarrow 1$, $\theta^i \leftarrow \theta_0^i$, $x \leftarrow x_0$
**for** $n = 1, \ldots, M$ **do**
Play action $a_n^i := \pi_n^i(x_n)$ along with other agents in current state $x_n \in \mathcal{X}$
Obtain next state $y_n \in \mathcal{X}$
Observe reward vector $r_n =< r_n^1, \ldots, r_n^N >$

**Value Update:** For $j = 1, \ldots, N$

$$v_{n+1}^j(x_n) = v_n^j(x_n) + c(n)\left(r_n^j + \beta v_n^j(y_n) - v_n^j(x_n)\right) \tag{13}$$

**Gradient Estimation:**

$$\xi_{n+1}^i(x_n, a_n^i) = \xi_n^i(x_n, a_n^i) + c(n)\Bigg(\sum_{j=1}^{N}\left(r_n^j + \beta v_n^j(y_n) - v_n^j(x_n)\right) - \xi_n^i(x_n, a_n^i)\Bigg) \tag{14}$$

**Policy Update:**

$$\pi_{n+1}^i(x_n, a_n^i) = \Gamma(\pi_n^i(x_n, a_n^i) - b(n)\sqrt{\pi_n^i(x_n, a_n^i)} \times \left|r_n^i + \beta v_n^i(y_n) - v_n^i(x_n)\right| \overline{\text{sgn}}(-\xi_{n+1}^i(x_n, a_n^i))) \tag{15}$$

**end for**

A few remarks about ON-SGSP are in order.

**Remark 6.** ***(Coupled dynamics)*** *In the ON-SGSP algorithm, an agent $i$ observes the rewards of other agents and uses this information to compute the respective value estimates. These quantities are then used to derive the decrement in the policy update* (15)*. This is meaningful in the light of the impossibility result of Hart and Mas-Colell [2003], where the authors show that in order to converge to a Nash equilibrium each agent's strategy needs to factor in the rewards of the other agents.*

**Remark 7.** ***(Per-iteration complexity)***

**OFF-SGSP:** *Let $A$ be the typical number of actions available to any agent in any given state and let $U$ be the typical number of next states for any state $x \in \mathcal{X}$. Then, the typical number of multiplications in OFF-SGSP per iteration is $N \times \left((U+1) \times A^N + 4A\right) \times |\mathcal{X}|$. Thus, the computational complexity grows exponentially in terms of the number of agents while being linear in the size of the state space. Note that the exponential behaviour in $N$ appears because of the computation of expectation over possible next states and strategies of agents. This computation is avoided in ON-SGSP.*

**ON-SGSP:** *For each agent, a typical iteration would take just $(2A+1)$ number of multiplications.*

*Thus, per-iteration complexity of OFF-SGSP is $\Theta(2^N)$ while that of ON-SGSP is $\Theta(1)$ (from the point of view of each agent). Thus, ON-SGSP is computationally efficient and this is also confirmed by simulation results, which establish that the total run-time till convergence of ON-SGSP is indeed very small when compared to that of the off-line algorithm OFF-SGSP. In comparison, the stochastic tracing procedure of Herings and Peeters [2004] has a complexity of $O(|\mathcal{X}| \times A^N)$ per iteration which is similar to that of OFF-SGSP.*

*However, the per-iteration complexity alone is not sufficient and an analysis of the number of iterations required is necessary to complete the picture*[8]*. On the other hand, convergence rate results for general multi-timescale stochastic approximation schemes are not available, see however, Konda and Tsitsiklis [2004] for rate results of two timescale schemes with linear recursions.*

**Remark 8.** ***(Descent directions)*** *It is shown in Proposition 1 that $-\sqrt{\pi^i(x,a^i)}\left|g^i_{x,a^i}(\mathbf{v}^i,\pi^{-i})\right| \operatorname{sgn}\left(\dfrac{\partial f(\mathbf{v},\pi)}{\partial \pi^i}\right)$ is a descent direction in $\pi^i(x,a^i)$ for every $i=1,2,\ldots,N,\ x \in \mathcal{X}, a^i \in \mathcal{A}^i(x)$. Since $\left\{\pi^i(x,a^i)\right\}^\alpha \geq 0$ for any $\alpha \geq 0$,*

$$-\left\{\pi^i(x,a^i)\right\}^\alpha \sqrt{\pi^i(x,a^i)}\left|g^i_{x,a^i}(\mathbf{v}^i,\pi^{-i})\right| \operatorname{sgn}\left(\frac{\partial f(\mathbf{v},\pi)}{\partial \pi^i}\right)$$

*can also be seen to be a descent direction in $\pi^i(x,a^i)$ for every $i=1,2,\ldots,N,\ x \in \mathcal{X}, a^i \in \mathcal{A}^i(x)$. In other words, the following is a descent direction in $\pi^i(x,a^i)$ for every $i=1,2,\ldots,N,\ x \in \mathcal{X}, a^i \in \mathcal{A}^i(x)$:*

$$-\left\{\pi^i(x,a^i)\right\}^{\alpha'}\left|g^i_{x,a^i}(\mathbf{v}^i,\pi^{-i})\right| \operatorname{sgn}\left(\frac{\partial f(\mathbf{v},\pi)}{\partial \pi^i}\right), \text{ for any } \alpha' \geq \frac{1}{2}$$

*So, the policy updates in OFF-SGSP/ON-SGSP can be generalized as follows: With $\alpha \geq \frac{1}{2}$,*

$$\textbf{OFF-SGSP:}\quad \pi^i(x,a^i) \leftarrow \Gamma\left(\pi^i(x,a^i) - \gamma\left\{\pi^i(x,a^i)\right\}^{\alpha'}\left|g^i_{x,a^i}(\mathbf{v}^i,\pi^{-i})\right| \overline{\operatorname{sgn}}\left(\frac{\partial f(\mathbf{v},\pi)}{\partial \pi^i}\right)\right),$$

$$\textbf{ON-SGSP:}\quad \pi^i(x,a^i) \leftarrow \Gamma\left(\pi^i(x,a^i) - \gamma\left\{\pi^i(x,a^i)\right\}^{\alpha'}\left|r^i + \beta v^i(y) - v^i(x)\right| \overline{\operatorname{sgn}}(q^i(x,a^i)\right),$$

*where $\gamma$ is a step-size parameter.*

**Remark 9.** ***(Convergence result)*** *Theorem 4 holds for ON-SGSP as well, while the proof deviates significantly. OFF-SGSP assumes model information, i.e., knowledge of transition dynamics. On the other hand, ON-SGSP operates in a model-free setting and hence, the analyses for both the timescales change. In particular, ON-SGSP uses a TD-critic and using standard stochastic approximation arguments (as in earlier literature), it is straightforward to prove the value updates in* (13) *converge to the true value function. However, the analysis of the $\xi$-recursion changes significantly. The latter is a consequence of the fact that ON-SGSP operates in a model-free setting and hence does not have access to $f(\mathbf{v}_n,\pi_n)$ (and hence $\dfrac{\partial f(\mathbf{v}_n,\pi_n)}{\partial \pi^i(x,a^i)}$ which is required for the policy update). Finally, the policy updates can be shown to track the system of ODEs* (19) *as in OFF-SGSP, after handling an additional martingale sequence that arises due to the $\xi$-recursion. The detailed proof is available in Section 8.2.*

[8] A well-known complexity result [Papadimitriou, 1994] establishes that finding the Nash equilibrium of a two-player game is PPAD-complete.

# 8 Proof of Convergence

We provide a proof of convergence of the two proposed algorithms - OFF-SGSP and ON-SGSP, respectively. In addition to Assumption 1, we make the following assumption for the analysis of our algorithms:

**Assumption 2.** *The underlying Markov chain with transition probabilities $p(y|x,\pi)$, $x, y \in \mathcal{X}$, corresponding to the general-sum discounted stochastic game, is irreducible and positive recurrent for all possible strategies $\pi$.*

The above assumption is standard in the analysis of multi-agent RL algorithms and can be seen in earlier works as well (for instance, see Hu and Wellman [1999], Littman [2001]).

In the following section, we provide the detailed analysis for OFF-SGSP and later, in Section 8.2, provide the necessary modifications to the analysis for ON-SGSP.

## 8.1 Proof of Theorem 4 for OFF-SGSP

As mentioned earlier, OFF-SGSP employs two time-scale stochastic approximation [Borkar, 2008, Chapter 6]. That is, it comprises of iteration sequences that are updated using two different time-scales or step-size schedules defined via $\{b(n)\}$ and $\{c(n)\}$, respectively. The step-sizes, satisfying Assumption 1, ensure the following:

**(i)** The policy $\pi$ (on slower timescale) appears quasi-static for updates of $\mathbf{v}$; and

**(ii)** The value $\mathbf{v}$ (on faster timescale) appears almost equilibrated for updates of $\pi$. We let $\mathbf{v}_\pi$ denote the value for a given policy $\pi$.

Claim (i) above can be inferred as follows: First rewrite the $\pi$-recursion in (9) as

$$\pi^i_{n+1}(x,a^i) = \Gamma\left(\pi^i_n(x,a^i) - c(n)H(n)\right),$$

where $H(n) = \frac{b(n)}{c(n)}\sqrt{\pi^i_n(x,a^i)}\left|g^i_{x,a^i}(\mathbf{v}^i_n, \pi_n^{-i})\right| \overline{\text{sgn}}\left(\frac{\partial f(\mathbf{v}_n,\pi_n)}{\partial \pi^i}\right)$, with $g^i(\cdot,\cdot)$ as defined in (4) and $f(\cdot,\cdot)$ as defined in (3). Since we consider a finite state-action spaced stochastic game, $g^i_{x,a^i}$ is bounded, while one can trivially upper bound $\pi$ and $\overline{\text{sgn}}$. Thus, $\sup_n |H(n)|$ is finite. Since, $\frac{b(n)}{c(n)} = o(1)$ by assumption 1, it can be clearly seen that the $\pi$-recursion in (9) tracks the ODE

$$\frac{d\pi^i(x,a^i)}{dt} = 0.$$

Claim (i) now follows.

Inferring claim (ii) above is technically more involved, but follows using arguments similar to those used in Theorem 2 of Chapter 6 in [Borkar, 2008].

In order to prove Theorem 4, we analyse each timescale separately in the following.

### Step 1: Analysis of $\mathbf{v}$-recursion

We first show that the updates of $\mathbf{v}$, that are on the faster time-scale, converge to a limit point of the following system of ODEs:$\forall x \in \mathcal{X}, i = 1, 2, \ldots, N$,

$$\frac{dv^i(x)}{dt} = r^i(x,\pi) + \beta \sum_{y \in U(x)} p(y|x,\pi)v^i(y) - v^i(x), \tag{16}$$

where $\pi$ is assumed to be time-invariant. We will also see that the system of ODEs above has a unique limit point, henceforth referred to as $\mathbf{v}_\pi$, which is stable.

Let $R^i_\pi = \left\langle r^i(x,\pi), x \in \mathcal{X}\right\rangle$ be a column vector of rewards to agent $i$ and $P_\pi = [p(y|x,\pi), x \in \mathcal{X}, y \in \mathcal{X}]$ be the transition probability matrix, both for a given $\pi$.

**Lemma 4.** *The system of ODEs (16) has a unique globally asymptotically stable limit point given by*

$$\mathbf{v}_\pi^i = [I - \beta P_\pi]^{-1} R_\pi^i, i = 1, 2, \ldots, N. \tag{17}$$

*Proof.* The system of ODEs (16) can be re-written in vector form as given below.

$$\frac{d\mathbf{v}^i}{dt} = R_\pi^i + \beta P_\pi \mathbf{v}^i - \mathbf{v}^i. \tag{18}$$

Rearranging terms, we get

$$\frac{d\mathbf{v}^i}{dt} = R_\pi^i + (\beta P_\pi - I)\mathbf{v}^i,$$

where $I$ is the identity matrix of suitable dimension. Note that for a fixed $\pi$, this ODE is linear in $\mathbf{v}^i$ with state transition matrix as $(\beta P_\pi - I)$. Since $P_\pi$ is a stochastic matrix, the magnitude of all its eigen-values is upper bounded by 1. Hence all the eigen-values of the state transition matrix $(\beta P_\pi - I)$ have negative real parts and the matrix $(\beta P_\pi - I)$ is in particular non-singular. Thus by standard linear systems theory, the above ODE has a unique globally asymptotically stable limit point which can be computed by setting $\frac{d\mathbf{v}^i}{dt} = 0, i = 1, 2, \ldots, N$, i.e.,

$$R_\pi^i + (\beta P_\pi - I)\mathbf{v}^i = 0, i = 1, 2, \ldots, N.$$

The trajectories of the ODE (18) converge to the above point starting from any initial condition in lieu of the above. □

For a given $\pi$, the updates of $\mathbf{v}$ in equation (10) (OFF-SGSP) can be seen as Euler discretization of the system of ODEs (16). We now show that $\mathbf{v}_n$ in equation (10) of OFF-SGSP converges to $\mathbf{v}_\pi$ as given in equation (17). While the following claim is identical for both OFF-SGSP/ON-SGSP, the proofs are quite different. In the former case, it amounts to proving value iteration converges (a standard result in dynamic programming), while the latter case amounts to proving a stochastic approximation variant of value iteration converges (also a standard result in RL).

**Proposition 2.** *For a given $\pi$, i.e., with $\pi_n^i \equiv \pi^i$, updates of $\mathbf{v}$ governed by (10) (OFF-SGSP) satisfy $\mathbf{v}_n \to \mathbf{v}_\pi$, as $n \to \infty$, where $\mathbf{v}_\pi$ is the globally asymptotically stable equilibrium point of the system of ODEs (16).*

*Proof.* We verify here assumptions (A1) and (A2) of Borkar and Meyn [2000] in order to use their result [Borkar and Meyn, 2000, Theorem 2.2]. Let $h(\mathbf{v}^i) = R_\pi^i + (\beta P_\pi - I)\mathbf{v}^i$. Since $h(\mathbf{v}^i)$ is linear in $\mathbf{v}^i$, it is Lipschitz continuous. Let $h_r(\mathbf{v}^i) = \frac{h(r\mathbf{v}^i)}{r}$ for a scalar real number, $r > 0$. It is easy to see that $h_r(\mathbf{v}^i) = \frac{R_\pi^i}{r} + (\beta P_\pi - I)\mathbf{v}^i$. Now, $h_\infty(\mathbf{v}^i) = \lim_{r\to\infty} h_r(\mathbf{v}^i) = (\beta P_\pi - I)\mathbf{v}^i$. Now since all eigenvalues of $(\beta P_\pi - I)$ have negative real parts, the ODE $\frac{d\mathbf{v}^i}{dt} = h_\infty(\mathbf{v}^i)$ has the origin as its unique globally asymptotically stable equilibrium. Further, as shown in Lemma 4, $\mathbf{v}_\pi^i$ is the unique globally asymptotically stable equilibrium for the ODE (18). Assumption (A1) of Borkar and Meyn [2000] is thus satisfied. Since the updates of $\mathbf{v}$ in equation (10) do not have any noise term in them, assumption (A2) of citet borkar2000ode is trivially satisfied. Thus by Borkar and Meyn [2000, Theorem 2.2], $\mathbf{v}_n$ in equation (10) converges to the globally asymptotically stable limit point $\mathbf{v}_\pi$ given in equation (17). □

Thus, on the faster time-scale $\{c(n)\}$, the updates of $\mathbf{v}$ obtained from (10) converge to $\mathbf{v}_\pi$, as given by (17).

## Step 2: Analysis of $\pi$-recursion

Using the converged values of $\mathbf{v}$ corresponding to strategy update $\pi_n$, i.e., $\mathbf{v}_{\pi_n}$ on the slower time-scale, we show that updates of $\pi$ converge to a limit point of the following system of ODEs:
For all $a^i \in \mathcal{A}^i(x), x \in \mathcal{X}, i = 1, 2, \ldots, N$,

$$\frac{d\pi^i(x, a^i)}{dt} = \bar{\Gamma}\left(-\sqrt{\pi^i(x, a^i)}\left|g_{x,a^i}^i(\mathbf{v}_\pi^i, \pi^{-i})\right| \overline{\text{sgn}}\left(\frac{\partial f(\mathbf{v}_\pi, \pi)}{\partial \pi^i}\right)\right), \tag{19}$$

where $g^i_{x,a^i}(\cdot,\cdot)$ is the Bellman error (see (4)), $f(\cdot,\cdot)$ is the objective in (3) and $\bar{\Gamma}$ is a projection operator that restricts the evolution of the above ODE to the simplex $\mathcal{D}$ and is defined as follows:

$$\bar{\Gamma}(v(x)) = \lim_{\eta \to 0} \left( \frac{\Gamma(x + \eta v(x)) - x}{\eta} \right), \tag{20}$$

for any continuous $v : \mathcal{D} \to \mathcal{R}^N$.

Let $K$ denote the limit set of the ODE (19). Before we analyse the $\pi$-recursion in (9), we show that the points in $K$ that are infeasible for the optimization problem (3) are asymptotically unstable. In other words, each stable limit point of the ODE (19) is an SG-SP point.

Define the set of all feasible points of the optimization problem (3) as follows:

$$G = \left\{ \pi \in L \middle| g^i_{x,a^i}(v^i_\pi, \pi^{-i}) \leq 0, \forall a^i \in \mathcal{A}^i(x), x \in \mathcal{X}, i = 1, 2, \ldots, N \right\} \tag{21}$$

The limit set $K$ of the ODE (19) can be partitioned using the feasible set $G$ as $K = K_1 \cup K_2$ where $K_1 = K \cap G$ and $K_2 = K \setminus K_1$. In the following lemma, we show that the set $K_2$ is the set of locally unstable equilibrium points of (19).

**Lemma 5.** *All $\pi^* \in K_2$ are unstable equilibrium points of the system of ODEs (19).*

*Proof.* For any $\pi^* \in K_2$, there exists some $a^i \in \mathcal{A}^i(x), x \in \mathcal{X}, i \in \{1, 2, \ldots, N\}$, such that $g^i_{x,a^i}(\mathbf{v}^i_\pi, \pi^{-i}) > 0$ and $\pi^i(x, a^i) = 0$ because $K_2$ is not in the feasible set $G$. Let $B_\delta(\pi^*) = \{\pi \in L | \, \|\pi - \pi^*\| < \delta\}$. Choose $\delta > 0$ such that $g^i_{x,a^i}(\mathbf{v}^i_\pi, \pi^{-i}) > 0$ for all $\pi \in B_\delta(\pi^*) \setminus K$ and consequently $\dfrac{\partial f(\mathbf{v}_\pi, \pi)}{\partial \pi^i} < 0$.

So, $\bar{\Gamma}\left( -\sqrt{\pi^i(x,a^i)} \left| g^i_{x,a^i}(\mathbf{v}^i_\pi, \pi^{-i}) \right| \operatorname{sgn}\left( \dfrac{\partial f(\mathbf{v}_\pi, \pi)}{\partial \pi^i} \right) \right) > 0$ for any $\pi \in B_\delta(\pi^*) \setminus K$ which suggests that $\pi^i(x, a^i)$ will increase when moving away from $\pi^*$. Thus, $\pi^*$ is an unstable equilibrium point of the system of ODEs (19). □

**Remark 10.** ***(On the sign function)*** *Recall that $\overline{\operatorname{sgn}}$ was employed since the normal $sgn()$ function is discontinuous. Since $\overline{\operatorname{sgn}}$ can result in the value $0$, one can no longer conclude that $\sqrt{\pi^*} g = 0$ for the points in the equilibrium set $K$. Note that the former condition (coupled with feasibility) implies it is an SG-SP point. A naive fix would be to change OFF-SGSP/ON-SGSP to repeat an action if $\overline{\operatorname{sgn}}(\cdot)$ returned $0$. This would ensure that there are no spurious points in the set $K$ due to $\overline{\operatorname{sgn}}$ being $0$. Henceforth, we shall assume that there are no such $\overline{\operatorname{sgn}}$ induced spurious limit points in the set $K$.*

**Lemma 6.** *For all $a^i \in \mathcal{A}^i(x), x \in \mathcal{X}$ and $i = 1, 2, \ldots, N$,*

$$\pi \in K \Rightarrow \pi \in L \text{ and } \sqrt{\pi^i(x,a^i)} g^i_{x,a^i}(v^i_\pi, \pi^{-i}) = 0, \tag{22}$$

*where $L = \left\{ \pi | \pi(x) \text{ is a probability vector over } \mathcal{A}^i(x), \forall x \in \mathcal{X} \right\}$.*

*Proof.* The operator $\bar{\Gamma}$, by definition, ensures that $\pi \in L$. Suppose for some $a^i \in \mathcal{A}^i(x), x \in \mathcal{X}$ and $i \in \{1, 2, \ldots, N\}$, we have $\bar{\Gamma}(-\sqrt{\pi^i(x,a^i)} \left| g^i_{x,a^i}(v^i_\pi, \pi^{-i}) \right| \overline{\operatorname{sgn}}\left( \dfrac{\partial f(\mathbf{v}_\pi, \pi)}{\partial \pi^i} \right)) = 0$, but $\sqrt{\pi^i(x,a^i)} g^i_{x,a^i}(v^i_\pi, \pi^{-i}) \neq 0$. Then, $g^i_{x,a^i}(v^i_\pi, \pi^{-i}) \neq 0$ and since $\pi \in L$, $1 \geq \pi^i(x, a^i) > 0$. We analyze this condition by considering the following two cases.

**Case $1 > \pi^i(x, a^i) > 0$ and $g^i_{x,a^i}(v^i_\pi, \pi^{-i}) \neq 0$.**

In this case, it is possible to find a $\Delta > 0$ such that for all $\delta \leq \Delta$,

$$1 > \pi^i(x, a^i) - \delta \sqrt{\pi^i(x,a^i)} \left| g^i_{x,a^i}(v^i_\pi, \pi^{-i}) \right| \overline{\operatorname{sgn}}\left( \frac{\partial f(\mathbf{v}_\pi, \pi)}{\partial \pi^i} \right) > 0.$$

This implies that

$$\begin{aligned}
&\bar{\Gamma}\left(-\sqrt{\pi^i(x,a^i)}\left|g^i_{x,a^i}(v^i_\pi,\pi^{-i})\right|\overline{\text{sgn}}\left(\frac{\partial f(\mathbf{v}_\pi,\pi)}{\partial \pi^i}\right)\right)\\
&=-\sqrt{\pi^i(x,a^i)}\left|g^i_{x,a^i}(v^i_\pi,\pi^{-i})\right|\overline{\text{sgn}}\left(\frac{\partial f(\mathbf{v}_\pi,\pi)}{\partial \pi^i}\right)\neq 0,\\
\Rightarrow&\sqrt{\pi^i(x,a^i)}g^i_{x,a^i}(v^i_\pi,\pi^{-i})\neq 0,
\end{aligned}$$

which contradicts the initial supposition.

**Case $\pi^i(x,a^i)=1$ and $g^i_{x,a^i}(v^i_\pi,\pi^{-i})\neq 0$.**

Since $\mathbf{v}^i_\pi$ is solution of the system of ODEs (16), the following should hold:

$$\sum_{\hat{a}^i\in\mathcal{A}^i(x)}\pi^i(x,\hat{a}^i)g^i_{x,\hat{a}^i}(v^i_\pi,\pi^{-i})=\pi^i(x,a^i)g^i_{x,a^i}(v^i_\pi,\pi^{-i})=0.$$

This again leads to a contradiction.

The result follows. □

In order to prove Theorem 4, we require the well-known Kushner-Clark lemma (see [Kushner and Clark, 1978, pp. 191-196]). For the sake of completeness, we recall this result below.

**Theorem 5.** ***(Kushner-Clark lemma)*** *Consider the following recursion in $d$-dimensions:*

$$x_{n+1}=\Gamma(x_n+b(n)(h(x_n)+\zeta_n+\beta_n)), \tag{23}$$

*where $\Gamma$ projects the iterate $x_n$ onto a compact and convex set, say $\mathcal{C}\in\mathrm{R}^d$. The ODE associated with (23) is given by*

$$\dot{x}(t)=\bar{\Gamma}(h(x(t))), \tag{24}$$

*where $\bar{\Gamma}$ is a projection operator that keeps the ODE evolution within the set $\mathcal{C}$ and is defined as in* (20)*. We make the following assumptions:*

***(B1)*** *$h$ is a continuous $\mathrm{R}^d$-valued function.*

***(B2)*** *The sequence $\beta_n, n\geq 0$ is a bounded random sequence with $\beta_n\to 0$ almost surely as $n\to\infty$.*

***(B3)*** *The step-sizes $b(n), n\geq 0$ satisfy $b(n)\to 0$ as $n\to\infty$ and $\sum_n b(n)=\infty$.*

***(B4)*** *$\{\zeta_n, n\geq 0\}$ is a sequence such that for any $\epsilon>0$,*

$$\lim_{n\to\infty}P\left(\sup_{m\geq n}\left\|\sum_{i=n}^{m}b(i)\zeta_i\right\|\geq\epsilon\right)=0.$$

*Suppose that the ODE* (24) *has a compact set $K^*$ as its set of asymptotically stable equilibrium points. Then, $x_n$ converges almost surely to $K^*$ as $n\to\infty$.*

## Proof of Theorem 4

*Proof.* The updates of $\pi$ given by (9) on the slower time-scale $\{b(n)\}$ can be rewritten as: For all $a^i\in\mathcal{A}^i(x), x\in\mathcal{X}$ and $i=1,2,\ldots,N$,

$$\pi^i_{n+1}(x,a^i)=\Gamma\left(\pi^i_n(x,a^i)-b(n)(H(\pi^i_n)+\beta_n)\right), \tag{25}$$

where

$$H(\pi_n^i) = \sqrt{\pi_n^i(x,a^i)}\left|g^i_{x,a^i}(v^i_{\pi_n},\pi_n^{-i})\right|\overline{\text{sgn}}\left(\frac{\partial f(v_{\pi_n},\pi_n)}{\partial \pi^i}\right),$$
$$\beta_n = \sqrt{\pi_n^i(x,a^i)}\left|g^i_{x,a^i}(v^i_{n},\pi_n^{-i})\right|\overline{\text{sgn}}\left(\frac{\partial f(v_{n},\pi_n)}{\partial \pi^i}\right) - H(\pi_n^i).$$

We now verify the assumptions (B1)–(B4) for the recursion above:

- $H(\pi_n^i)$ is continuous since each of its components $\sqrt{\pi^i}$, $\left|g^i_{x,a^i}(\cdot,\cdot)\right|$ and $\overline{\text{sgn}}(\cdot)$ are continuous. In particular, the continuity of $g^i_{x,a^i}$ follows from the fact that both the value function $\mathbf{v}^i(\cdot)$ and Q-value function $Q^i_{\pi^{-i}}(x,a^i)$ are continuous in $\pi^i$. This verifies assumption (B1).
- $\beta_n \to 0$ almost surely since $|v_n - v_{\pi_n}| \to 0$ as $n \to \infty$, from Theorem 2. Further, $\beta_n$ is bounded as each of its components are bounded. In particular, $g^i_{x,a^i}$ is bounded as we consider finite state-action spaced stochastic games, while $\pi^i$ and $\overline{\text{sgn}}$ are trivially upper-bounded. Thus (B2) is satisfied.
- Assumption 1 implies (B3) is satisfied.
- $\zeta_n$ is absent, obviating assumption (B4).

The claim now follows from Kushner-Clark lemma. □

**Remark 11.** *Note that from the foregoing, the set $K$ comprises of both stable and unstable attractors and in principle from Lemma 5, the iterates $\pi_n^i$ governed by (19) can converge to an unstable equilibrium. In most practical scenarios, however, a gradient descent scheme is observed to converge to a stable equilibrium. In fact, the $\delta$-offset policy computed using $perturb(\cdot,\delta)$ for every $Q > 0$ iterations (see Section 9 below) for both of our algorithms ensures numerically that as $n \to \infty$, $\pi_n \nrightarrow \pi^* \in K_2$. In other words, convergence of the strategy sequence $\pi_n$ governed by* (9) *is to the stable set $K_1$.*

## 8.2 Proof of Theorem 4 for ON-SGSP

As mentioned earlier, the analysis for ON-SGSP changes for both timescales and we outline the crucial differences below, before presenting the detailed analysis.

**Step 1:** This step establishes that the TD updates along faster timescale converge to the true value functions, using standard techniques from stochastic approximation. Unlike OFF-SGSP, this step also involves the analysis of the $\xi$-recursion. The latter is a consequence of the fact that we work in a model-free setting and hence do not have access to $f(\mathbf{v}_n,\pi_n)$ (and hence $\dfrac{\partial f(\mathbf{v}_n,\pi_n)}{\partial \pi^i(x,a^i)}$ which is required for the policy update).

**Step 2:** This step establishes that the policy updates track the same ODE (i.e., (19)) as that of OFF-SGSP and the analysis involves an additional martingale sequence that needs to be bounded.

### Step 1: Analysis of $\mathbf{v}$ and $\xi$-recursions

**Proposition 3.** *For a given $\pi$, i.e., with $\pi_n^i \equiv \pi^i$, updates of $\mathbf{v}$ governed by (13) (ON-SGSP) satisfy $\mathbf{v}_n \to \mathbf{v}_\pi$ almost surely as $n \to \infty$, where $\mathbf{v}_\pi$ is the globally asymptotically stable equilibrium point of the system of ODEs (16).*

*Proof.* Fix a state $x \in \mathcal{X}$. Let $\{\bar{n}\}$ represent a sub-sequence of iterations in ON-SGSP when the state is $x \in \mathcal{X}$. Also, let $Q_n = \{\bar{n} : \bar{n} < n\}$. For a given $\pi$, the updates of $\mathbf{v}$ on the faster time-scale $\{c(n)\}$ given in equation (13) can be re-written as

$$v^i_{\bar{n}+1}(x) = v^i_{\bar{n}}(x) + c(\bar{n})\left[J(v^i_{\bar{n}}) + \tilde{\chi}_{\bar{n}}\right],$$

where

$$J(v^i_{\bar{n}}) = \sum_{a^i \in \mathcal{A}^i(x)} \pi^i(x,a^i) g^i_{x,a^i}(\mathbf{v}^i_{\bar{n}}, \pi^{-i}), \text{ and}$$

$$\tilde{\chi}_{\bar{n}} = \big(r^i_n + \beta v^i_n(y_n) - v^i_n(x_n)\big) - \sum_{a^i \in \mathcal{A}^i(x)} \pi^i(x,a^i) g^i_{x,a^i}(\mathbf{v}^i_{\bar{n}}, \pi^{-i}).$$

Using arguments as before, it is easy to see that $J(v^i_{\bar{n}})$ is continuous, $\tilde{\chi}_{\bar{n}}$ is such that $\mathbb{E}\tilde{\chi}^2_{\bar{n}} < \hat{C} < \infty$. Thus,

$$\lim_{\bar{n}\to\infty} P\left(\sup_{m \geq \bar{n}} \left\| \sum_{l=\bar{n}}^{m} c(l)\tilde{\chi}_l \right\| \geq \epsilon \right) \leq \frac{\hat{C}}{\epsilon^2} \lim_{\bar{n}\to\infty} \sum_{l=\bar{n}}^{\infty} c(l)^2 = 0.$$

For the last equality, we have used the fact that the step-sizes are square-summable (see assumption 1). Thus. all the assumptions of Kushner-Clark Lemma (see Theorem 5 above) are satisfied and we can conclude that $v_n$ governed by (13) converges to the globally asymptotically stable limit point $\mathbf{v}_\pi$ (see equation (17)) of the system of ODEs (16). □

Before establishing the convergence of the gradient estimation recursion, i.e., (14), we require the following technical result.

**Lemma 7.**

$$\frac{\partial f(\mathbf{v},\pi)}{\partial \pi^i(x,a^i)} = -\sum_{j=1}^{N} g^j_{x,a^i}(v^j, \pi^{-i}), \text{ where} \tag{26}$$

$$g^j_{x,a^i}(\mathbf{v}^j, \pi^{-i}) = \bar{r}^j(x,\pi^{-i},a^i) + \beta \sum_{y \in U(x)} \bar{p}^j(y|x,\pi^{-i},a^i) v^j(y) - v^j(x). \tag{27}$$

*Proof.* Let $a \in \mathcal{A}(x)$ denote the aggregate action vector. Then, we have

$$\begin{aligned}
\frac{\partial f(\mathbf{v},\pi)}{\partial \pi^i(x,a^i)} =& -g^i_{x,a^i}(v^i,\pi^{-i}) - \sum_{j\neq i} \sum_{a^j \in \mathcal{A}^j(x)} \pi^j(x,a^j) \Bigg( \sum_{k\neq i,j} \sum_{a^k \in \mathcal{A}^k(x)} \prod_{k\neq i,j} \pi^k(x,a^k) \big( r^j(x,a) \\
&+ \beta \sum_{y\in U(x)} p(y|x,a) v^j(y) - v^j(x) \big) \Bigg) \\
=& -g^i_{x,a^i}(v^i,\pi^{-i}) - \sum_{j\neq i} \Bigg( \sum_{k\neq i} \sum_{a^k \in \mathcal{A}^k(x)} \prod_{k\neq i} \pi^k(x,a^k) \big( r^j(x,a) \\
&+ \beta \sum_{y\in U(x)} p(y|x,a) v^j(y) - v^j(x) \big) \Bigg).
\end{aligned}$$

Now, from the definition of $g^j_{x,a^i}(v^i,\pi^{-i})$ in (27), it is easy to see that

$$\frac{\partial f(\mathbf{v},\pi)}{\partial \pi^i(x,a^i)} = -\sum_{j=1}^{N} g^j_{x,a^i}(v^j, r^j, \pi^{-i}).$$

□

Recall that the gradient estimation recursion is as follows:

$$\xi^i_{n+1}(x_n,a^i_n) = \xi^i_n(x_n,a^i_n) + c(n)\Bigg( \sum_{j=1}^{N} \big( r^j_n + \beta v^j_n(y_n) - v^j_n(x_n) \big) - \xi^i_n(x_n,a^i_n) \Bigg).$$

The following theorem establishes that $\xi^i_n(x,a^i)$ converges to $-\dfrac{\partial f(\mathbf{v}_n,\pi_n)}{\partial \pi^i(x,a^i)}$ in the long run.

**Proposition 4.** $\left\| \xi_n^i(x,a^i) - \left( -\frac{\partial f(\mathbf{v}_n,\pi_n)}{\partial \pi^i(x,a^i)} \right) \right\| \to 0$ *as* $n \to \infty$ *almost surely.*

*Proof.* As in the proof of Proposition 3, let $\{\bar{n}\}$ represent a sub-sequence of iterations in ON-SGSP when the state is $x \in \mathcal{X}$ and $Q_n := \{\bar{n} : \bar{n} < n\}$. For a given $\pi$, the updates of $\xi$ on the faster time-scale $\{c(n)\}$ given in equation (14) can be re-written as

$$\xi^i_{\bar{n}+1}(x,a^i) = (1-c(\bar{n}))\xi^i_{\bar{n}}(x,a^i) + c(\bar{n})\Big( \sum_{j=1}^{N} g^j_{x,a^i}(v^j_{\bar{n}}, r^j_{\bar{n}}, \pi^{-i}) + \xi_{\bar{n}} \Big),$$

where $\xi_{\bar{n}} := \sum_{j=1}^{N} \left( r^j_{\bar{n}} + \beta v^j_{\bar{n}}(y) - v^j_{\bar{n}}(x) \right) - \sum_{j=1}^{N} g^j_{x,a^i}(v^j_{\bar{n}}, r^j_{\bar{n}}, \pi^{-i})$. Let $\mathcal{F}_l := \sigma(\mathbf{v}_k, \xi_k, k \le l), l \ge 0$ denote an increasing family of $\sigma$-fields. By definition of $g^j_{x,a^i}$, we have

$$E\left[ r^j_{\bar{n}} + \beta v^j_{\bar{n}}(y) - v^j_{\bar{n}}(x) \mid \mathcal{F}_{\bar{n}}, \pi^i_{\bar{n}}(x_{\bar{n}}, a^i_{\bar{n}}) \right] = \sum_{j=1}^{N} g^j_{x,a^i}(v^j_{\bar{n}}, r^j_{\bar{n}}, \pi^{-i}).$$

Hence, $\{\xi_{\bar{n}}\}$ is a martingale difference sequence.

As in Proposition 3, define $\tilde{M}_n = \sum_{m \in Q_n} c(m)\tilde{\chi}_m$. It can be easily verified that $(\tilde{M}_m, \mathcal{F}_m), m \ge 0$ is a square-integrable martingale sequence obtained from the corresponding martingale difference $\{\xi_m\}$. Further, from the square summability of $c(n), n \ge 0$, and assumption 2 which ensures that the underlying Markov chain is ergodic for any given $\pi$, it can be verified from the martingale convergence theorem that $\{\tilde{M}_m, m \ge 0\}$, converges almost surely. Hence, $|\xi_m| \to 0$ almost surely on the 'natural timescale', as $m \to \infty$. The 'natural timescale' is clearly faster than the algorithm's timescale and hence $\xi_m$ can be ignored in the analysis of Bellman error recursion ((14) in the main paper), see [Borkar, 2008, Chapter 6.2] for detailed treatment of natural timescale algorithms. The final claim follows from Kushner-Clark lemma. □

### Step 2: Analysis of $\pi$-recursion

*Proof.* We first re-write the update of $\pi$ as follows: For all $i = 1, 2, \ldots, N$,

$$\pi^i_{\bar{n}+1}(x,a^i) = \Gamma\left( \pi^i_{\bar{n}}(x,a^i) - b(\bar{n}) \left( \sqrt{\pi^i_{\bar{n}}(x,a^i)} \left| g^i_{x,a^i}(v^i_{\pi_{\bar{n}}}, \pi^{-i}_{\bar{n}}) \right| \overline{\text{sgn}}\left( \frac{\partial f(v_{\pi_{\bar{n}}}, \pi_{\bar{n}})}{\partial \pi^i} \right) + \zeta_{\bar{n}} \right) \right),$$

where $\zeta_{\bar{n}} = \sqrt{\pi^i_{\bar{n}}(x,a^i)} \left[ \left| \hat{g}^i_{x,a^i} \right| - \left| g^i_{x,a^i}(v^i_{\pi_{\bar{n}}}, \pi^{-i}_{\bar{n}}) \right| \right] \overline{\text{sgn}}\left( \frac{\partial f(v_{\pi_{\bar{n}}}, \pi_{\bar{n}})}{\partial \pi^i} \right)$. In the above, we have used the converged values of the value update $v_n$ and gradient estimate $\xi_n$ and this is allowed due to timescale separation and the fact that $v_n \to \mathbf{v}_{\pi_n}$ (see Proposition 3) and $\xi_n \to -\frac{\partial f(\mathbf{v}_n,\pi_n)}{\partial \pi^i(x,a^i)}$ (see Proposition 4). Now, in order to apply Kushner-Clark lemma (see Theorem 5 above), it is enough if we verify that $\mathbb{E}\zeta^2_{\bar{n}} < \infty$, since the rest of the terms are as in OFF-SGSP (which imply assumptions (B1) to (B3) in Theorem 5 are verified). Now, arguing as before, it is straightforward to infer that $\mathbb{E}\zeta^2_{\bar{n}} < \infty$, since we consider finite state-action spaces and the square of each of the quantities in the first term in $\zeta_{\bar{n}}$ can be upper-bounded. Thus, assumption (B4) in Theorem 5 is verified and updates of $\pi$ in ON-SGSP converge to a stable limit point of the system of ODEs (19). □

# 9 Simulation Experiments

We test ON-SGSP, NashQ [Hu and Wellman, 2003] and FFQ [Littman, 2001] algorithms on two general-sum game setups. We implemented Friend Q-learning variant of FFQ, as each iteration of its Foe Q-learning variant involves a computationally intensive operation to solve a linear program.

| **Player 2 → Player 1 ↓** | $a_1$ | $a_2$ | $a_3$ |
|---|---|---|---|
| $a_1$ | $1, 0$ | $0, 1$ | $1, 0$ |
| $a_2$ | $0, 1$ | $1, 0$ | $1, 0$ |
| $a_3$ | $0, 1$ | $0, 1$ | $1, 1$ |

(a) Payoff matrix.

| | **NashQ** | **FFQ (Friend Q)** | **ON-SGSP** |
|---|---|---|---|
| **Oscillate or converge to non-Nash strategy** | **95**% | **40**% | 0% |
| **Converge to** $(0.5, 0.5, 0)$ | 2% | 0% | **99**% |
| **Converge to** $(0, 0, 1)$ | 3% | 60% | 1% |

(b) Results from 100 simulation runs.

Figure 4: Payoff matrix and simulation results for a single state non-generic two-player game

## 9.1 Single State (Non-Generic) Game

This is a simple two-player game adopted from Hart and Mas-Colell [2005], where the payoffs to the individual agents are given in Table 4a. In this game, a strategy that picks $a_3$ (denoted by $(0, 0, 1)$) constitutes a pure-strategy NE, while a strategy that picks either $a_1$ or $a_2$ with equal probability (denoted by $(0.5, 0.5, 0)$) is a mixed-strategy NE.

We conduct a stochastic game experiment where at each stage, the payoffs to the agents are according to Table 4a and the payoffs accumulate with a discount factor $\beta = 0.8$. We performed $100$ experimental runs, with each run corresponding to a length of $10000$ stages. The aggregated results from this experiment are presented in Fig. 4b. It is evident that NashQ oscillates and does not converge to NE in most of the runs, while Friend Q-learning converges to a non-Nash strategy tuple in most of the runs. On the other hand, ON-SGSP converges to NE in all the iterations.

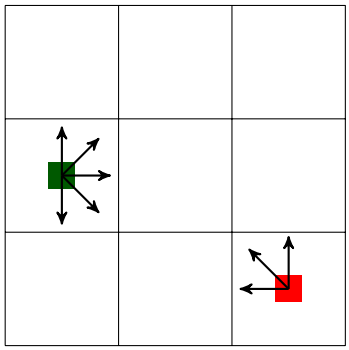

Figure 5: Stick-Together Game for $M = 3$

## 9.2 Stick-Together Game (STG)

We also define a simple general-sum discounted stochastic game, named "Stick-Together Game" or in short STG, where two participating agents located on a rectangular terrain would like to come together and stay close to each other (see Fig. 5). A precise description of the various components of STG is provided below:

1. **State Space $\mathcal{X}$:** The state specifies the location of both the agents on a rectangular grid of size $M \times M$. More precisely, let $O = \{(x, y)|x, y \in \mathbb{Z}\}$. Denote the possible positions of an agent by $W := \{s = (x, y) \in O|0 \leq x, y < M\}$. Then the state space is given by the Cartesian product $\mathcal{X} = W \times W$.

2. **Action Space $\mathcal{A}$:** The actions available to each agent are to either move to one of the neighboring cells or stay in the current location. For $s \in W$, let $\|s\|_1 = |x| + |y|$ be its $L^1$ norm. Then, $A(s) = \{a \in O|\|s + a\| \leq 1\}$ represents the actions available for an agent to move to one-step neighbouring positions of $s \in W$. The action space is then defined by $\mathcal{A} = \underset{s^1, s^2 \in W}{\cup} A(s^1) \times A(s^2)$. Let $U(s) = \{s' \in W|\|s' - s\|_1 \leq 1\}$ represent the set of all next states for an agent in state $s \in W$.

3. **Transition probability $p$:** We assume that state transitions of individual agents are independent. Let $q(s'|s, a^i)$ represent, for agent $i$, the probability of transition from state $s \in W$ to $s' \in W$ upon taking action $a^i \in \mathcal{A}^i(s)$. We define

$$q(s'|s, a) = \frac{2^{-\|s'-a\|_1}}{\sum\limits_{s'' \in U(s)} 2^{-\|s''-a\|_1}}.$$

   Then, the transition probability is given by

$$p((s'^1, s'^2)|(s^1, s^2), (a^1, a^2)) = q(s'^1|s^1, a^1)q(s'^2|s^2, a^2).$$

   This transition probability function has the highest value towards that next state to which the action points to.

4. **Reward $r$:** The reward for the two agents is defined as

$$r^i(s^i, a^i) = 1 - e^{\|s^1 - s^2\|_1},$$

   for state $(s^1, s^2) \in \mathcal{X}$ and action $(a^1, a^2) \in A(s^1) \times A(s^2)$. Thus, the reward is zero if the distance between the two agents is zero. Otherwise, it is a negative and monotonically decreasing function with respect to the distance between the two agents.

**Results.** We first show simulation results for a small sized version of the STG game, where $M = 3$. The number of states with $M = 3$ is $|\mathcal{X}| = 81$. We use $\beta = 0.8$ for all our experiments. Also, we use the following step-size sequences in our experiments:

$$b(n) = \begin{cases} 0.2 & \text{for } n < 1000, \\ \dfrac{1}{n^{0.75}} & \text{otherwise,} \end{cases}$$

$$c(n) = \begin{cases} 0.1 & \text{for } n < 1000, \\ \dfrac{1}{n} & \text{otherwise.} \end{cases}$$

It is easy to see that $\{c(n)\}$ corresponds to slower time-scale than $\{b(n)\}$. It was observed in our experiments that, using constant step-sizes upto $n = 1000$ leads to better initial exploration and faster convergence.

To ensure sufficient exploration of the state space (see assumption 2) and also to push the policy $\pi$ out of the domain of attraction of any local equilibrium, we perturb the policy as follows: For every $Q > 0$ iterations, $perturb(\cdot, \delta)$ is used to derive a $\delta$-offset policy for picking actions, i.e., $\hat{\pi}^i(x)$ is used instead of $\pi^i(x)$, where

$$\hat{\pi}^i(x, a^i) = \frac{\pi^i(x, a^i) + \delta}{\sum\limits_{a^i \in \mathcal{A}^i(x)} (\pi^i(x, a^i) + \delta)}, a^i \in \mathcal{A}^i(x). \tag{28}$$

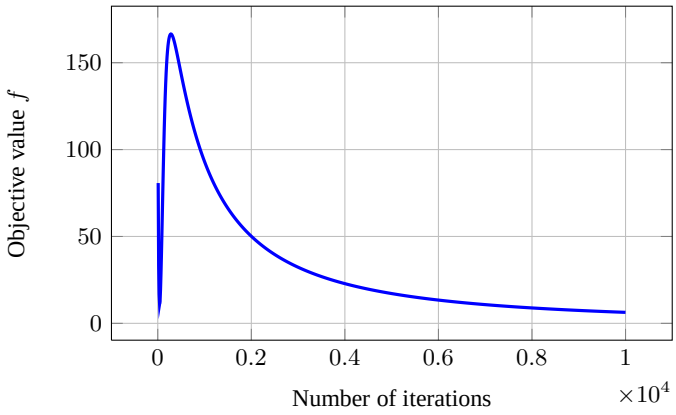


(a) OFF-SGSP for STG with $M = 3$

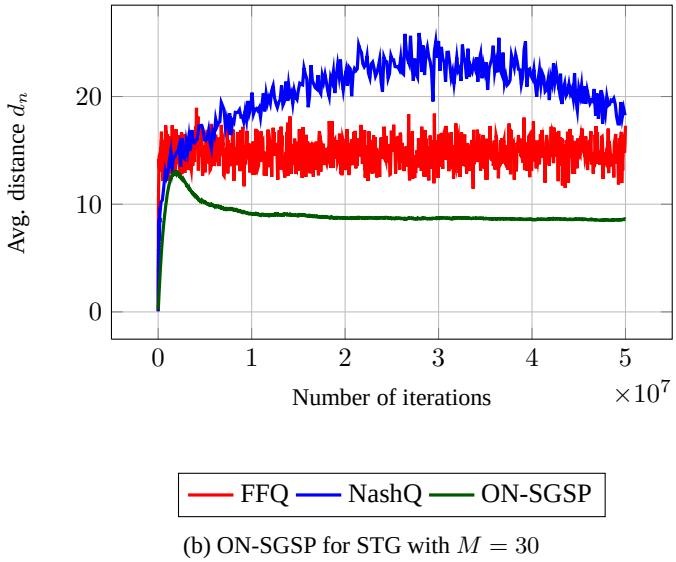


(b) ON-SGSP for STG with $M = 30$

Figure 6: Performance of our algorithms for STG

Fig. 6a shows the evolution of the objective function $f$ as a function of the number of iterations for OFF-SGSP. Note that $f$ should go to zero for a Nash equilibrium point.

Fig. 6b shows the evolution of the distance $d_n$ (in $\ell_1$ norm) between the agents for a STG game where $M = 30$, which corresponds to $|\mathcal{X}| = 810,000$. Notice that the results are shown only for the model-free algorithms: ON-SGSP, NashQ and FFQ. This is because OFF-SGSP and even the homotopy methods [Herings and Peeters, 2004] have exponential blow up with $M$ in their computational complexity and hence, are practically infeasible for STG with $M = 30$.

From Fig. 6b, it is evident that following the ON-SGSP strategy, the agents converge to a $4 \times 4$-grid within the $30 \times 30$-grid. For achieving this result, ON-SGSP takes about $2 \times 10^7$ iterations, implying an average $2 \times 10^7/|\mathcal{X}| \approx 21$ iterations per state. However, NashQ gets the agents to an $8 \times 8$-grid after a large number of iterations ($\approx 5 \times 10^7$). Moreover, from Fig. 6b it is clear that NashQ has not stabilized its strategy in the end. Friend Q-learning gets the agents to $8 \times 8$-grid, by driving them to one of the corners of the $30 \times 30$-grid. While it takes a short number of iterations ($\approx 30000$) to achieve this, FFQ does not explore the state space well and hence, FFQ's strategy corresponding to the rest of the grid (excluding the corner to which it takes the agents) is not Nash.

**Remark 12.** ***(Runtime performance.)*** *We observed that to complete $5 \times 10^7$ iterations, ON-SGSP took $\approx 42$*

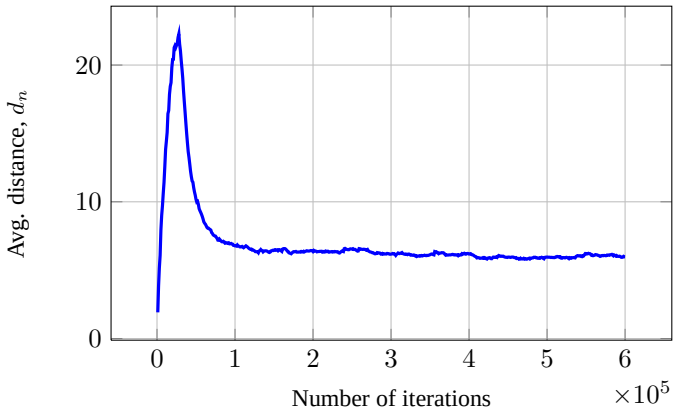


Figure 7: ON-SGSP with partial information for STG with $M = 30$

*minutes, while NashQ [Hu and Wellman, 2003] took nearly* $50$ *hours, as it involves solving for Nash equilibria of a bimatrix game in each iteration. The Friend Q-learning variant of FFQ [Littman, 2001] took* $\approx 33$ *minutes. The Foe Q-learning variant of FFQ was not implemented owing to its high per-iteration complexity.*

### Simple Function Approximation for STG

While OFF-SGSP assumes full information of the game, ON-SGSP assumes that neither rewards nor state transition probabilities are known. Here, we explore an intermediate information case albeit restricted to STG where a partial structure of rewards is made known. In particular, we assume that the reward depends on the difference $\Delta = (|x_1^1 - x_1^2|, |x_2^1 - x_2^2|) \in \mathcal{X}$ in positions $x_1 = (x_1^1, x_2^1), x_2 = (x_1^2, x_2^2) \in \mathcal{X}$ of the two agents. We approximate the value function $\mathbf{v}$ and strategy $\pi$ as follows: $v^i(x) \approx \hat{v}^i(\Delta)$ and $\pi^i(x) \approx \hat{\pi}^i(\Delta), \forall x \in \mathcal{X}$. Thus, the algorithms need to compute $\hat{\mathbf{v}}$ and $\hat{\pi}$ on a low-dimensional subspace $W$ of $\mathcal{X}$. Fig. 7 presents results of ON-SGSP in this setting for $M = 30$, which corresponds to $|W| = 900$. The solution is seen to have converged by $200,000$ iterations ($\approx 5$ seconds runtime) which suggests that it took on an average $\frac{200,000}{|W|} \approx 22$ iterations per $\Delta \in W$ to converge.

## 10 Conclusions

In this paper, we derived necessary and sufficient SG-SP conditions to solve a generalized optimization problem and established their equivalence with Nash strategies. We derived a descent (not necessarily steepest) direction that avoids local minima. Incorporating this direction, we proposed two algorithms - offline, model-based algorithm OFF-SGSP and online, model-free algorithm ON-SGSP. Both algorithms were shown to converge, in self-play, to the equilibria of a certain ordinary differential equation (ODE), whose stable limit points coincide with stationary Nash equilibria of the underlying general-sum stochastic game. Synthetic experiments on two general-sum game setups show that ON-SGSP outperforms two well-known multi-agent RL algorithms. The experimental evaluation also suggests that convergence is relatively quick.

There are several future directions to be explored and we outline a few of them below:

1. In simulations, we observed that ON-SGSP can successfully run for large state spaces ($|\mathcal{X}| \approx 8,00,000$). However, in many cases, the state spaces can be huge and it would be necessary to look for function approximation techniques for both the value function $\mathbf{v}$ as well as strategy-tuple $\pi$. Function approximation techniques are popular in reinforcement learning approaches for high-dimensional MDPs and they bring in the following advantages: (a) They can cater to huge state and action spaces; and (b) They also aid a designer to bring in his understanding about the underlying system in terms of features used for function approximations.

2. Extensions to the case of constrained games: By constrained stochastic games, we mean those stochastic games that have additional constraints on value functions or strategy-tuple which might arise from the modelling of a practical scenario. There are some results and applications in this by Altman et al. [2005] and Altman et al. [2007] which provide the necessary motivation for extending our results to general-sum constrained stochastic games.

3. Detailed experimental evaluation on a sophisticated benchmark for $N$-player general-sum stochastic games.

# References


N. Akchurina. Multi-agent Reinforcement learning: Algorithm converging to Nash equilibrium in general-sum stochastic games. In *8th International Conference on Autonomous Agents and Multi-agent Systems*, 2009.

E. Altman, K. Avrachenkov, R. Marquez, and G. Miller. Zero-sum constrained stochastic games with independent state processes. *Mathematical Methods of Operations Research*, 62(3):375–386, 2005.

E. Altman, K. Avratchenkov, N. Bonneau, M. Debbah, R. El-Azouzi, and D.S. Menasche. Constrained stochastic games in wireless networks. In *IEEE Global Telecommunications Conference*, pages 315–320. IEEE, 2007.

V.S. Borkar. *Stochastic Approximation: A Dynamical Systems Viewpoint*. Cambridge Univ Press, 2008.

V.S. Borkar and S.P. Meyn. The ODE method for convergence of stochastic approximation and reinforcement learning. *SIAM Journal on Control and Optimization*, 38(2):447–469, 2000.

R. N. Borkovsky, U. Doraszelski, and Y. Kryukov. A user's guide to solving dynamic stochastic games using the homotopy method. *Operations Research*, 58(4-Part-2):1116–1132, 2010.

M. Bowling. Convergence and no-regret in multiagent learning. *Advances in neural information processing systems*, 17:209–216, 2005.

M. Bowling and M. Veloso. Rational and convergent learning in stochastic games. In *International joint conference on artificial intelligence*, volume 17, pages 1021–1026, 2001.

M. Breton. *Algorithms for stochastic games*. Springer, 1991.

V. Conitzer and T. Sandholm. Awesome: A general multiagent learning algorithm that converges in self-play and learns a best response against stationary opponents. *Machine Learning*, 67(1-2):23–43, 2007.

J. Filar and K. Vrieze. *Competitive Markov Decision Processes*. Springer-Verlag, New York, Inc., 1st edition, November 2004. ISBN 0-387-94805-8.

A. M. Fink. Equilibrium in a stochastic $n$-person game. *Journal of Science of the Hiroshima University, Series AI (Mathematics)*, 28(1):89–93, 1964.

S. Hart and A. Mas-Colell. Uncoupled dynamics do not lead to Nash equilibrium. *American Economic Review*, pages 1830–1836, 2003.

Sergiu Hart and Andreu Mas-Colell. Stochastic uncoupled dynamics and nash equilibrium. In *Proceedings of the 10th conference on Theoretical aspects of rationality and knowledge*, pages 52–61. National University of Singapore, 2005.

P. J. Herings and R. Peeters. Homotopy methods to compute equilibria in game theory. Research Memoranda 046, Maastricht : METEOR, Maastricht Research School of Economics of Technology and Organization, 2006. URL `http://ideas.repec.org/p/dgr/umamet/2006046.html`.

P. J. Herings and R. J. A. P. Peeters. Stationary equilibria in stochastic games: Structure, selection, and computation. *Journal of Economic Theory*, 118(1):32–60, 2004.

J. Hu and M. P. Wellman. Multiagent reinforcement learning: Theoretical framework and an algorithm. In *Proceedings of 15th International Conference on Machine Learning*, pages 242–250, 1999.

J. Hu and M. P. Wellman. Nash Q-Learning for general-sum stochastic games. In *Journal of Machine Learning Research*, volume 4, pages 1039–1069, 2003.

E. Kalai and E. Lehrer. Rational learning leads to Nash equilibrium. *Econometrica: Journal of the Econometric Society*, pages 1019–1045, 1993.

V. R. Konda and J. N. Tsitsiklis. Convergence rate of linear two-time-scale stochastic approximation. *Annals of Applied Probability*, pages 796–819, 2004.

H.J. Kushner and D.S. Clark. *Stochastic approximation methods for constrained and unconstrained systems*. Springer-Verlag, 1978. ISBN 0-387-90341-0.

M. L. Littman. Markov games as a framework for multi-agent reinforcement learning. *ICML*, 94:157–163, 1994.

M. L. Littman. Friend-or-Foe Q-Learning in General Sum Games. In *Proceedings of the 18th International Conference on Machine Learning*, pages 322–328. Morgan Kaufmann, 2001.

L. Mac Dermed and C. L. Isbell. Solving stochastic games. *Advances in Neural Information Processing Systems*, 22:1186–1194, 2009.

C. H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *Journal of Computer and System Sciences*, 48(3):498–532, 1994.

H. L. Prasad and S. Bhatnagar. General-sum stochastic games: verifiability conditions for Nash equilibria. *Automatica*, 48(11):2923–2930, nov 2012.

H. L. Prasad and S. Bhatnagar. A Study of Gradient Descent Schemes for General-Sum Stochastic Games. *arXiv preprint arXiv:1507.00093*, 2015.

H. L. Prasad, L.A. Prashanth, and S. Bhatnagar. Two-timescale algorithms for learning nash equilibria in general-sum stochastic games. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pages 1371–1379, 2015.

L. S. Shapley. Stochastic games. In *Proceedings of the National Academy of Sciences*, volume 39, pages 1095–1100, 1953.

Y. Shoham, R. Powers, and T. Grenager. Multi-agent reinforcement learning: a critical survey. *Web manuscript*, 2003.

Y. Shoham, R. Powers, and T. Grenager. If multi-agent learning is the answer, what is the question? *Artificial Intelligence*, 171(7):365–377, 2007.

M. J. Sobel. Noncooperative stochastic games. *The Annals of Mathematical Statistics*, 42(6):1930–1935, 1971.

R. S. Sutton. Learning to predict by the methods of temporal differences. *Machine learning*, 3(1):9–44, 1988.

R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA, 1998. ISBN 0-262-19398-1.

M. Takahashi. Equilibrium points of stochastic non-cooperative $n$-person games. *Journal of Science of the Hiroshima University, Series AI (Mathematics)*, 28(1):95–99, 1964.

J. W. Weibull. *Evolutionary Game Theory*. MIT Press, 1996.

C. Zhang and V. R. Lesser. Multi-agent learning with policy prediction. In *AAAI*, 2010.

M. Zinkevich, A. Greenwald, and M. Littman. Cyclic equilibria in Markov games. *Advances in Neural Information Processing Systems*, 18:1641–1648, 2006.